Intel Developer Update is Intel's monthly online news magazine for developers. As the official publication of developer.intel.com, it brings hardware, software, and Web developers the latest information on Intel initiatives, technologies, platforms, and products.

Cover Story

Each month, we run a cover story on the most significant industry announcement, trend, or development for the month.

Featured Articles

Delivering in-depth reports on key platforms, products and technologies, our featured articles provide a monthly source of information on issues affecting developers. Be sure to check in every month for the latest developments driving the evolution of the industry.

Contact the Editor

To make *Intel Developer Update* a better information resource, we invite you to share your thoughts on what we've published or what you'd like to see covered. Comments are always welcome.

Archives

Our archives contain two groups of previously published articles. One group contains all the articles that appeared in *Platform Solutions News*, the earlier version of *Intel Developer Update*. The articles date from September 1997 through August 1999. The other group is set up to contain *Intel Developer Update* articles dating from the inaugural September/October 1999 issue.

Bookmarking

We advise against bookmarking article pages. They're accessible online only during the month the issue is live. Thereafter, they're removed to our archives. Instead, we suggest that you bookmark the PDF (Adobe* Portable Document Format) file versions of the articles. You'll find buttons for the PDF files labeled "print article" in the right navigation section of each article. A PDF for the entire issue is labeled "print magazine" and is located near top right side of the IDU home page.

Table of Contents

(Click on page number to jump to articles)

Note: Intel does not control the content on other company's Web sites or endorse other companies supplying products or services. Any links that take you off of Intel's Web site are provided for your convenience.

Cover Story

Revolutionizing Enterprise Building Blocks: Intel® Xeon™ Processor-based Servers with the Intel® E7500 Chipset

Ryan C. Hebeler
Platform Applications Engineer
Enterprise Platforms Group
Intel Corporation

Overview

The Intel® E7500 chipset-based platform represents Intel's next-generation volume server chipset and is the first of several enterprise-class chipsets targeted for introduction in 2002. With this platform comes the simultaneous launch of three new products plus the availability of high-performance network connectivity provided by PCI-X Gigabit Ethernet.

- Intel® Xeon™ processor with 512-KB L2 Cache**
- Intel® E7500 Chipset
- Intel® IOP321 I/O processor
- Intel® 82544GC and Intel® 82546EB Gigabit Ethernet Controllers

**Note: Unless otherwise stated in this article, "Intel® Xeon™ processor" refers to the Intel® Xeon™ processor with 512-KB L2 Cache.

The high-throughput Intel E7500 Chipset-based Dual-Processor (DP) server platform takes full advantage of the Intel Xeon processor to deliver unparalleled compute power, value, and versatility for Internet infrastructure and departmental/SMB (Small and Medium Business) server applications.

The block diagram shown in Figure 1.0 illustrates how the Intel Xeon processor and the Intel E7500 Memory Controller Hub (MCH) connect to the various interfaces to provide a complete platform solution for DP server applications. The system bus interface and the memory interface as well as the hub interface 2.0 connections to the I/O subsystem depict the maximum available bandwidth of 3.2 GB/s each. The ICH3-S attaches to the Intel E7500 MCH through a 266-MB/s hub interface 1.5 connection.

Figure 2.0, in the "High-Performance I/O with 64-bit PCI/PCI-X Technology" section of this article, illustrates how devices such as I/O processors and Gigabit Ethernet controllers can attach to 64-bit PCI-PCI-X Controller Hub 2 (P64H2) devices through the PCI-X connections.

This article provides developers and integrators with a high-level understanding of the features and advantages offered by this platform. For more detailed information on the products, features, and interfaces of the Intel E7500 Chipset-based DP server platform, see the "More Info" section of this article.

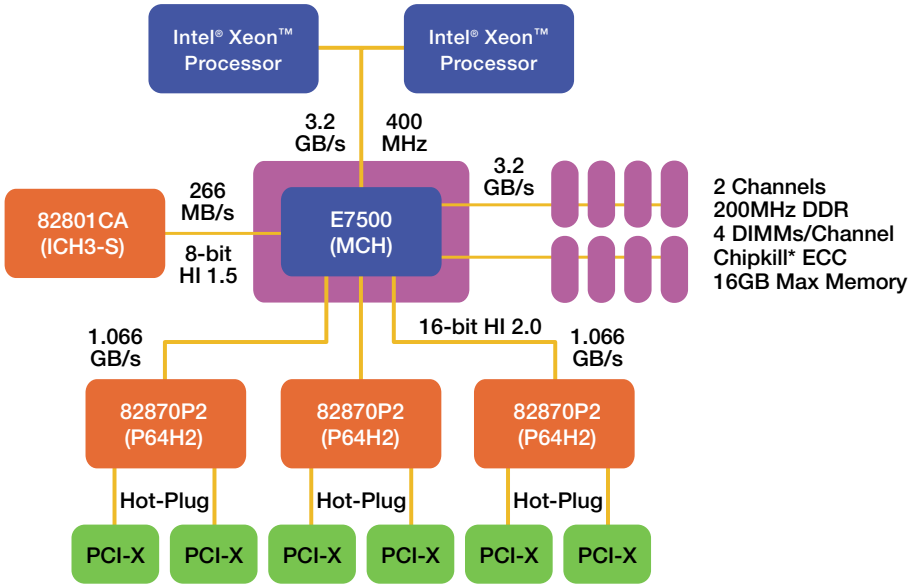

Figure 1. Intel® Xeon™ Processor and the Intel® E7500 Chipset Dual Processor Server Block Diagram

Introducing Intel® Xeon™ Processor with 512-KB L2 Cache

The Intel Xeon processor with 512-KB L2 Cache is the first Intel processor for DP servers to include the powerful Intel® NetBurst™ microarchitecture, the first new 32-bit microarchitecture in five years.

This new microarchitecture takes advantage of a 400-MHz system bus to deliver higher throughput when accessing memory and I/O devices to deliver improved server headroom and scalability. A Rapid Execution Engine utilizes 2X clock speeds for integer computations, increasing performance for Web and database servers. Together, the higher throughput and the faster clock times result in a hyper-pipelined technology that enables higher transaction rates and faster response times. In addition to these benefits, the availability of 144 new Streaming SIMD Extensions 2 (SSE-2) instructions will improve response times for media servers, secure transactions, and next-generation Web services.

The Intel NetBurst microarchitecture is just one of several new features of the Intel Xeon processor. This microarchitecture provides benefits across a broad range of server applications, and it works together with other new features such as Hyper-Threading Technology and an Integrated Cache Subsystem to make the Intel Xeon processor the ideal choice for a DP server environment.

Hyper-Threading Technology Plus Integrated Cache Equals Performance

Along with The Intel NetBurst microarchitecture, Hyper-Threading Technology and an Integrated Cache Subsystem also make their debut in the enterprise market segment.

The Intel NetBurst microarchitecture and Hyper-Threading Technology work together to deliver best-in-class performance for peak Internet server workloads. Enhancements in bandwidth, throughput, and thread-level parallelism deliver an acceleration of performance for applications such as search engines, security, and streaming media. The integrated cache subsystem is tightly coupled with the Intel NetBurst microarchitecture to offer greater throughput for today's large server workloads. These advancements in technology and performance have yielded the headroom to support an ever-increasing number of users in an expanding e-Business environment.

From a software standpoint, Hyper-Threading Technology is groundbreaking technology that provides immediate, tangible benefits for today's multi-threaded server applications. New software is not required, as it uses existing multi-threading capability in operating systems and applications. This technology allows software to utilize processor resources that would otherwise be idle, increasing performance up to 30 percent when used with today's multi-threaded server software.

This performance increase has a significant, noticeable impact on the performance of server applications:

- Increases the number of transactions that can be processed for the enterprise
- Enables support for more users, improving business productivity
- Provides faster response times for many Internet and e-Business applications, enhancing the customer's experience

Looking forward, Hyper-Threading Technology will enable new capabilities and applications for the evolving Internet and enterprise infrastructure in areas such as Web services, personalization, clustering, and consolidation. The industry has already begun development focused on the optimization of operating systems and applications to utilize the capabilities of this new processor and further the benefits of threading. Intel has worked with various operating system vendors on products such as Windows* XP, Windows* .NET Server, and Linux* to develop and include optimizations for Hyper-Threading Technology in future releases of their products.

Intel also continues to enable middleware and application tuning for the Intel NetBurst microarchitecture and Hyper-Threading Technology. Technical support for these applications, plus development tools such as Intel® Compilers, Performance Libraries, and the VTune Performance Analyzer, are available.

While Hyper-Threading Technology complements multiprocessing by providing additional headroom for future software optimizations and business growth, an integrated cache subsystem utilizes several features to meet the needs of today's server applications. Super-fast Level-1 (L1) Data Cache and Execution Trace Cache are timed and synchronized with the NetBurst pipeline and the Rapid Execution engine to reduce access latency and maximize pipeline throughput. A 512-KB L2 Advanced Transfer Cache is tightly synchronized with the L1 Data cache and the Rapid Execution engine to improve access times.

Peak available bandwidth of 70.0 GB/s for this cache subsystem more than triples the available bandwidth of previous generation cache architectures (2.2-GHz Intel Xeon processor being compared to a 1.4-GHz Pentium® III Xeon™ processor with a total available bandwidth of 22.4 GB/s). This increase in available cache bandwidth provides a much better match for the increase in bandwidth required by the Intel Xeon processor execution units.

Satisfying the demand of these execution units as they consume data is the key to maximizing throughput of the Intel Xeon processor core. Improved cache hit rates are made possible by increased cache line sizes (128-byte lines compared to 32-byte lines for Pentium III Xeon processor).

All these performance enhancements introduced by the integrated cache subsystem of the Intel Xeon processor add up to improved overall performance for today's and tomorrow's server applications.

Platform-Level Features and Performance Enhancements

At a platform level, Intel Xeon processor-based servers introduce significant improvements in system bus performance along with important advancements in memory and I/O technologies. The result is new levels of platform bandwidth and throughput for enterprise server applications.

The Intel E7500 chipset has been optimized for the Intel Xeon processor to yield the scalability and headroom necessary for future business growth. At 3.2 GB/s, the system bus interface provides a 3X increase in system bus bandwidth over previous IA-32-based servers. This processor and chipset combination provides best-in-class, balanced platform performance for demanding server workloads, and it enables a smooth transition to next-generation server technologies.

The availability of this platform also marks the high-volume introduction of Double Data Rate (DDR) memory technology to the enterprise server market space. This technology allows the completion of two data operations in one clock cycle—twice the throughput of regular SDRAM—yielding increased headroom for memory-intensive applications and faster data mining. The 200-MHz DDR memory technology (DDR-200) delivers both faster data transfer and a 1.5X increase in memory bandwidth over yesterday's PC-133, SDRAM-based DP server systems.

Along with the performance improvements offered by the system bus interface and the memory interface, the Intel E7500 chipset offers several high-bandwidth I/O configuration options that improve the I/O performance, flexibility, and expandability for a new generation of volume DP servers.

This Intel E7500 chipset-based DP server platform is the first Intel® server platform to include PCI-X I/O. The 64-bit, 133-MHz PCI-X technology doubles the frequency at which data is exchanged between the 64-bit PCI-PCI-X Controller Hub 2 (P64H2) and peripherals when compared to previous platforms. This performance increase can be utilized by high-bandwidth devices such as Gigabit Ethernet cards, Fiber Channel, and Ultra3 SCSI devices.
At 3.2 GB/s total I/O bandwidth, this I/O subsystem offers a 3X increase in peak I/O bandwidth over the Intel® Pentium® III and Intel Pentium III Xeon family processor-based platforms.

Original Equipment Manufacturers (OEMs) and server application developers can benefit from the complete platform solution and the measurable performance improvements provided in the Intel E7500 chipset-based DP server platform. The advantages offered by this new platform can help improve existing server applications and product lines and help define new ones. Significant platform feature enhancements, performance advancements and raw value make this the time to upgrade from older solutions. Intel Xeon processor-based DP Servers can become an integral part of a business's competitive advantage.

Intel® E7500 MCH Delivers Balanced Platform Performance

The Intel E7500 chipset Memory Controller Hub, the MCH, manages the flow of data between the Dual Intel Xeon Processors and the other platform components. By matching the available throughput of the system bus, the memory, and the I/O interfaces at 3.2 GB/s each, the MCH maximizes the efficiency of data flow, creating a balanced performance environment for the platform. This performance-balanced platform approach takes full advantage of the available bandwidth for all platform interfaces, and offers exceptional reliability and scalability for a DP server system across a broad range of DP server applications.
The performance improvements brought by the Intel E7500 MCH can be summarized with the relative total bandwidth improvements over today's IA-32-based servers:

- Up to 3X increase in system bus bandwidth
- Up to 1.5X increase in memory bandwidth
- Up to 3X increase in I/O bandwidth

These platform-level performance improvements translate to measurable performance improvements in DP server applications: 25 to 30 percent more orders processed, requests serviced, and users supported than with today's Pentium III processor-based DP servers. Increased reliability means less system downtime while increased scalability provides the headroom necessary for future business growth.

High-Performance I/O with 64-bit PCI/PCI-X Technology

The Intel® 82870P2 64-bit PCI-PCI-X Controller Hub 2 (P64H2) is an I/O Bridge that performs PCI bridging functions between the hub interface and the PCI bus. The Intel E7500 MCH supports up to three P64H2 devices via its three hub interface 2.0 connections.

At over 1 GB/s each, these three hub interface connections combine to yield the 3.2 GB/s total bandwidth for the entire I/O subsystem for the platform. Once again, this is a 3X increase in peak I/O bandwidth over the Intel Pentium III and Intel Pentium III Xeon family processor-based platforms.

On the primary bus, the P64H2 utilizes a 16-bit data bus to attach to the hub interface 2.0 connection of the Intel E7500 MCH. The hub interface Error Correcting Code and the high throughput offered by the hub interface 2.0 connection combine to deliver both greater reliability and faster access to the high-speed networks for the platform.

On the secondary bus of the P64H2, there are two independent 64-bit, 133-MHz PCI-X interfaces, each with its own dedicated hot plug controller to enhance serviceability and its own I/OxAPIC with 24 interrupts. These secondary PCI

bus interfaces can be configured independently to operate in either PCI or PCI-X mode. These features allow I/O flexibility, including platform configuration options with up to six high-performance 64-bit, 133-MHz PCI-X buses.

The introduction of next-generation 64-bit PCI/PCI-X performance also allows high-throughput connectivity to Intel® Gigabit Ethernet and Intel® I/O processors. Figure 2.0 illustrates how Intel Gigabit Ethernet controllers (Intel® 82544GC and Intel® 82546EB) as well as the I/O processor (Intel® IOP321) can attach to the P64H2 to provide yet another layer of I/O configuration flexibility.

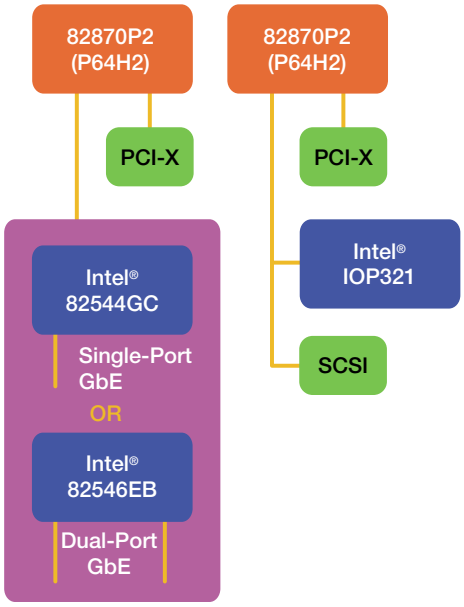

Figure 2. Intel® 82544GC or the Intel® 82546EB Gigabit Ethernet Controllers and the Intel® IOP321 I/O Processor — Flexible I/O Configuration Options for High-Speed Network Capability.

On-board Gigabit Ethernet: Intel® 82544GC and Intel® 82546EB Gigabit Ethernet Controllers

The Intel 82544GC Single-port Gigabit Ethernet controller matches the performance of the Intel E7500 chipset-based platform. This single-port solution provides the high-throughput network connections required by today's servers by incorporating a 133-MHz, 64-bit PCI-X interface that matches the I/O bandwidth of the P64H2 component. Integration of the MAC and PHY for 1000BASE-T Gigabit Ethernet offers the exceptional performance of the Intel Gigabit Ethernet products in a single, integrated component. The small size and low power consumption of the 82544GC make it ideal for Gigabit Ethernet LAN on motherboard (LOM) platform implementations.

In addition to the single-port Gigabit Ethernet solution, Intel recently announced the Intel 82546EB Dual Port Gigabit Ethernet Controller for the Intel E7500 Chipset-based platform. This dual-port controller represents only one load on the PCI-X bus, allowing it to operate at 133 MHz. At this frequency, it can support bandwidth to allow wire-speed performance of two Gigabit Ethernet connections. Whether decisions are being made at a board-design level, a system-design level, or at the marketing/product line level, the increase in performance and flexibility this dual-port controller offers over a single-port solution, coupled with its impressive feature set, make it the right choice:

- A motherboard design engineer with limited space and cooling requirements could choose the Intel 82546EB Dual Port Gigabit Ethernet Controller as a dual-port LOM solution to achieve performance gains and additional flexibility over a single-port solution without having to worry about a significant additional BOM cost.
- A system design engineer needing more I/O bandwidth for a new back-end server design can choose the Intel 82546EB Gigabit Ethernet controller and gain the added I/O throughput offered by two Gigabit Ethernet ports and PCI-X.

Whether single-port or dual-port controllers are chosen, Intel Gigabit Ethernet products are cost-effective solutions that offer the performance, flexibility, and advanced RASUM features required by server motherboard designers, server system integrators, and IT server management and staff.

I/O Processing with the Intel® IOP321 I/O Processor

The Intel IOP321 I/O processor is Intel's fifth-generation I/O processor. It is the first I/O processor to integrate an Intel® XScale™ microarchitecture core and a PCI-X interface. The IOP321 is a highly integrated, cost-effective I/O system on a chip that delivers a two-fold performance boost over its predecessor, the IOP310 I/O processor chipset, in I/O-intensive applications.

The Intel IOP321 I/O processor is especially well suited for networked storage applications including RAID (Redundant Array of Independent Disks) adapter cards, ROMB (RAID on motherboard), and other storage applications. This combination of features and advancements makes the Intel IOP321 an ideal choice for applications requiring a high-performance I/O subsystem in a tightly integrated environment.

This I/O processor brings added value and performance to the Intel E7500 chipset by integrating a memory controller unit and a RAID 5 accelerator. The integration of the RAID 5 acceleration unit allows an SCSI U320 disk controller to reside on the same PCI-X segment as the IOP321 with the RAID stack running on the IOP321. The integration of this task into the IOP321 returns available bandwidth to the platform, as CPU bandwidth is no longer being consumed for running RAID 5. This combination provides high-performance RAID acceleration at a low cost.

Legacy Support with the Intel® ICH3-S

The Intel® 82801CA I/O Controller Hub 3, the ICH3-S, connects to the Intel E7500 chipset MCH through a point-to-point, parity-protected hub interface 1.5 connection with an available bandwidth of 266 MB/s. The ICH3-S provides legacy I/O interfaces through integrated features such as a two-channel Ultra ATA/100 bus master IDE controller, three USB controllers for up to six total USB 1.0 ports and a 32-bit PCI interface. The ICH3-S also offers an integrated System Manageability Bus controller and an integrated LAN controller as well as AC'97 2.2-compliant and PCI 2.2-compliant interfaces.

The System Manageability Bus Controller, together with the other features of the ICH3-S, delivers the RASUM and legacy components that result in a reliable and compatible server platform for the enterprise.

Advanced RASUM Features

The Intel E7500 chipset-based server platform sets new expectations for the dual processor enterprise system market with its advanced RASUM features. These features complement exceptional platform performance with the reliability required in an enterprise server:

- Memory Error Correcting Code with Chipkill*
- Hardware memory scrubbing
- Hub interface Error Correcting Code
- Advanced out-of-band manageability
- Enhanced error status information maintained through reset

Summary

The features provided by both the Intel Xeon processor for DP Servers and the Intel E7500 chipset enable the delivery of stable technology for today's e-Business solutions. An innovative chipset design that optimizes I/O and memory bandwidth, together with the revolutionary Intel NetBurst microarchitecture, optimize performance, scalability, and end-user productivity, while providing a solid base for future technology.

Intel's building block infrastructure ensures fast deployment of next-generation enterprise platforms and technologies by incorporating key silicon components to the platform:

- Intel Xeon processor with 512-KB L2 Cache and Intel NetBurst Microarchitecture
- Intel E7500 Chipset
- Intel 82544GC Gigabit Ethernet Controller
- Intel IOP321 I/O processor

Intel's commitment to best-in-class validation and platform support make Intel Xeon processor-based DP server platforms with the new Intel E7500 Chipset the ideal choice for servers today and in the future.

More Info

General Information

Developer Web site
Other Intel Support

General Information Hotline: 800.628.8686/916.356.3104, 5 am–5 pm PST
Intel Literature Center: 800.548.4725, 7 am–7 pm CST
International Locations: Please contact your local sales office

Intel® Xeon™ Processor with 512-KB L2 Cache and Intel® NetBurst™ microarchitecture

Hyper-Threading Technology

Intel® E7500 Chipset Memory Controller Hub (MCH)

- Intel Doc # 298649—001: Intel® Xeon™ Processor with 512-KB L2 Cache and Intel® E7500 Chipset Platform Design Guide
- Intel Doc # 298647—001: Intel® E7500 Chipset Design Guide: E7500 Memory Controller Hub (MCH) Thermal and Mechanical Design Guidelines
- Intel Doc # 290730—001: Intel® E7500 Chipset Datasheet: E7500 Memory Controller Hub (MCH)
- Intel Doc # 290731—001: Intel® E7500 Chipset Memory Controller Hub (MCH) Specification Update

Intel® chipsets home pages

Intel® 82870P2 64-bit PCI-PCI-X Controller Hub 2 (P64H2)

- Intel Doc # 290732—001: Intel® 82870P2 PCI/PCI-X 64-bit Hub 2 (P64H2) Datasheet
- Intel Doc # 290735—001: Intel® 82870P2 PCI/PCI-X 64-bit Hub 2 (P64H2) Specification Update

Intel® 82801CA I/O Controller Hub 3-S (ICH3-S)

- Intel Doc # 290733—001: Intel® 82801CA I/O Controller Hub 3 (ICH3-S) Datasheet
- Intel Doc # 290739—001: Intel® 82801CA I/O Controller Hub 3 (ICH3-S) Specification Update

Intel® IOP321 I/O Processor

- Intel Doc # 273525: Intel® IOP321 I/O Processor Product Brief
- Intel Doc # 273517: Intel® 80321 I/O Processor Developer's Manual
- Intel Doc # 273520: Intel® 80321 I/O Processor Design Guide
- Intel Doc # 273518: Intel® 80321 I/O Processor Datasheet
- Intel Doc # 273519: Intel® 80321 I/O Processor Specification Update
- Intel Doc # 273691: Intel® IQ80321 Developer Kit Product Brief
- Intel Doc # 273521: Intel® IQ80321 Evaluation Board Manual

Application Notes

- Intel Doc # 273570: Intel® 80321 Design Review Checklist
- Intel Doc # 273522: Intel® 80321 Initialization Document

I/O Processor Page
Bridges Home Page

Intel® 82544GC and Intel® 82546EB Gigabit Ethernet Controllers

- Intel Doc # AP-427: 82544GC Gigabit Ethernet Controller Datasheet and Hardware Design Guide Application Note

Intel® Gbit Ethernet Controllers

Author Bio

Ryan Hebeler is a senior platform applications engineer in Intel's Enterprise Platforms Group. His nine-year Intel career also includes work in the Americas Application Support Group, America's Quality Support Center, the Folsom Design Center (writing architecture validation and silicon validation code), and Intel Platform Support Labs (performing signal quality testing and power delivery analysis). Ryan holds a bachelor's degree in electrical engineering from the University of Michigan.

Departments

Applied Computing

Ultra-Dense IA-Based PrPMCs Reduce Design Time and Cost

Mark Summers
Technical Marketing Engineer
Embedded Intel Architecture Division
Intel Corporation

Overview

Overcoming form factor, low power, thermal and computing requirements to create powerful, affordable embedded applications has resulted in the development of an Intel® Architecture-based processor peripheral component interconnect (PCI) mezzanine card (PrPMC). PrPMC technology drastically increases available real estate in small form factor designs. The approach used builds on new standard specifications and delivers the developer advantages of the robust Intel Architecture development ecosystem. Using Intel Architecture PrPMC technology, system designers can now utilize ultra-compute density while substantially reducing time-to-market and simplifying design updates.

Background

In the past, designers of embedded solutions facing power and form factor constraints developed proprietary solutions because standards governing the technologies that met their design requirements didn't exist. The development of proprietary solutions required substantial engineering effort and resulted in high-priced single-purpose elements with few common or reusable parts.

Density, which promises greater computing power in smaller and smaller footprints, has long been a technology trend of interest to these designers. Increasingly important in applications such as Web and media services, router and switch controls and other similar applications, density is often measured in terms of processors per rack, based on a standard rack comprising 42 U vertical rack units. For instance filling a rack with 1 U dual processor chassis would provide 84 processors per rack. Using Intel® density optimized servers, which contain eight servers in a 2 U chassis, such a rack would hold 168 processors. Ultra-density, as the term is used in this article, is density well beyond this level—the equivalent of more than 300 processors in a 42 U standard rack.

The promise of ultra-density is that substantially greater processing power becomes available in smaller, denser configurations, more likely to meet the form factor requirements that systems designers contend with. In addition, the reduced real estate required on the printed circuit board greatly simplifies the board design, resulting in lower development cost and quicker time-to-market.

Driven by pressure to reduce design costs, standards are emerging that support interoperability and interchangeability in embedded applications. The benefit of broadly accepted standards is commonality. With commonality comes greater choice and flexibility for designers and lower overall prices. Organizations such as the VMEbus International Trade Association (VITA), PCI Industrial Computer Manufacturing Group (PICMG), PCI Special Interest Group (PCISIG), and other organizations were formed to champion industry standard specifications.

Intel® Architecture in Ultra-Dense Applications

Intel Architecture was once considered too thermally hot to support ultra-dense applications. While fighting the war to keep heat and voltage requirements down, Intel incorporated new features and improvements made possible by advances in semiconductor technology to produce Intel Architecture processors ideal for ultra-dense applications. Figure 1 summarizes the thermal and other key properties of Intel® processors targeted for ultra-dense applications.

Processor	Core Speed	L2 Cache	Bus Speed (MHz)	Thermal Design Power	Core Vcc	Temp Case	Package Type	Silicon Tech
Celeron® (300) Ultra-Low Power	300MHz	128K	100	5.7W	1.1V	0-100° C	BGA2	0.18µm
Celeron® (400A) Low Power	400MHz	128K	100	10.1W	1.35V	0-100° C	BGA2	0.18µm
Pentium® III (400) Low Power	400MHz	256K	100	10.1W	1.35V	0-100° C	BGA2	0.18µm
Pentium® III (500) Low Power	500MHz	256K	100	12.2W	1.35V	0-100° C	BGA2	0.18µm
Pentium® III (800) Low Power	800MHz	512K	133	11.2W	1.15V	0-100° C	uFCBGA	0.13µm

Figure 1. Intel Embedded Processor Specifications

The Intel® Pentium® III processor is fully software compatible with previous members of the Intel® microprocessor family. Its history in desktop computing has resulted in a legacy of development and application software transferable to embedded applications. In conjunction with existing personal computer hardware and off-the-shelf development packages, Intel Pentium processor technology has become the ultimate embedded computing solution.

In short, the benefits of applying Intel Architecture to ultra-dense applications include:

- Robust development tools and developer community
- Scalable technology path
- Choice of input/output (I/O) buses
- Support from a broad range of existing software

The Case for PrPMC

Traditionally, computing system developers design a baseboard, which contains the computing engine parts. The baseboard contains interconnections for various specialized peripheral boards (or cards) for functions such as I/O. As technology evolved, however, designers found that the computing engine technology on the baseboard became obsolete faster than the specialized I/O card technologies.

The PrPMC was developed to solve this problem flipping the design equation around. The PrPMC puts the “quick to obsolete” computing engine parts (which are the processor, memory, chipset, and voltage regulator) on an add-in card and places the longer life span, specialized I/O functions on the baseboard. This approach allows the baseboard design to be dramatically simpler and smaller, lowering design time and time-to-market. When the processor becomes obsolete, upgrading involves merely unplugging the outdated card and replacing it with a new one, reducing the time and complexity of upgrading. Together, these lower total cost of ownership and provide excellent scalability.

Figure 2 shows a standard baseboard compared to an equivalent redesigned, ultra-dense PrPMC system.

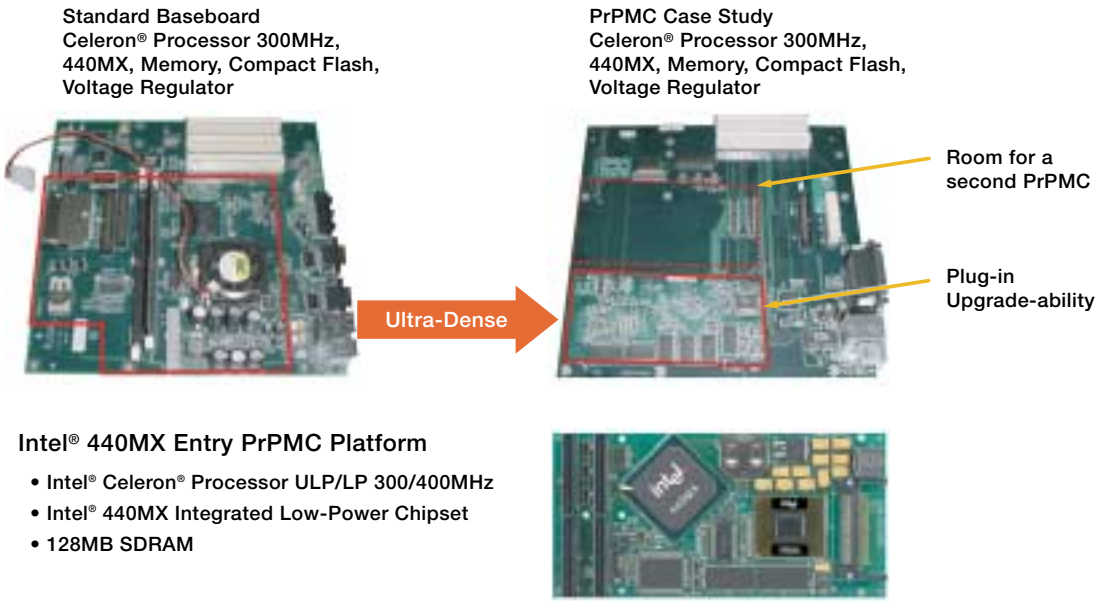

Figure 2. The advantages of ultra-density are substantially smaller board real estate for a design, greatly reducing design and development time. All functions of the standard baseboard on the left can be performed by the design on the right, which applies the ultra-dense design methodology described in this article. The photo in the lower right corner shows a close up of the Intel architecture based PrPMC, in this case, one of Intel design.

Case Study

The remainder of this article reviews a PrPMC design created by RadiSys Corporation. The design goal was to replace a standard baseboard design that utilizes an Intel® Pentium® III processor with an Intel Pentium III mobile processor-based PrPMC design. Figure 3 shows a block diagram of the RadiSys design.

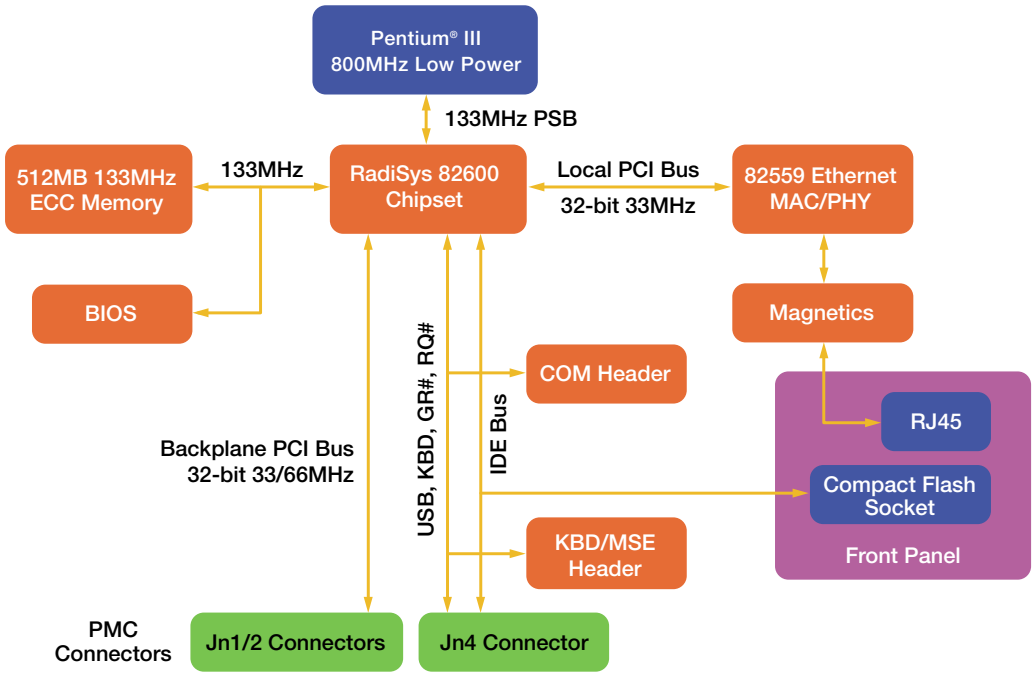

Figure 3. Block Diagram of the RadiSys EPC®-6315 PrPMC (Pentium® III Mobile Processor, RadiSys 82600 Chipset, 512MB ECC SDRAM).

Building an entire compute engine (processor, chipset, memory, and voltage regulator) on a PMC-sized (74 mm x 149 mm) printed circuit board is a daunting challenge. Further, the design constraints of integrating the PrPMC into a system need to be considered quite carefully due to interactions on the PCI bus, power allocation, and available system cooling. The design considerations RadiSys faced were:

- Layout, such as real estate (will it fit?) and routing and timing (will it work?)
- System integration
- Thermal characteristics

Figure 4 lists these design issues in greater detail.

Design Considerations	Description	
Printed Circuit Board Layout	The printed circuit board layout met the following requirements: • Overall thickness of 0.063 ±0.007 inches • Multiple potential planes (3.3V, VCORE, 5V, and GND) • 6 routing layers needed = 14 layers • Buried capacitance used • Components on both sides of the board • IEEE1386 envelope	
System Integration	Bus Arbitration	The host/non-host considerations for PCI bus (cPCI = Host, PrPMC = Non-Host): • The Non-Host PrPMC BIOS runs and operating system boots before the Host enumerates the bus
	System Cooling	System Cooling Requirements: • Approximately 30W/cPCI blade • 200 LFM in cPCI slot • Shared with all components • Host processor (or PrPMC) preheats the air and can block much of the flow to the Non-Host PrPMC • Maintain greater than 100 LFM over each PrPMC
	Other Form Factors	Non-blocking with different numbers of PrPMCs installed
	Power Allocation	Power allocation considerations with the processor on a cPCI card: • 3.3V is fairly heavily loaded • Needed to allocate early, mid and late power to host and PrPMC to meet loading specifications
Thermal Considerations	Thermal heat design issues included: • Z-height restrictions • Available surface area • Air flow rate • Ambient air temperature • Heat-generating neighbors	

Figure 4. Design considerations used in the RadiSys Intel PrPMC design.

Summary

Intel Architecture can be driven into ultra-density platform components, as demonstrated in this RadiSys PrPMC case study. The benefits include exceptional flexibility in selecting from the wide selection of robust tools optimized, ported or tuned for Intel Architecture to shorten development cycle time; a scalable product design path to maximize return on engineering investment and reduced product costs.

More Info

The design collateral you need to begin to apply this design breakthrough to your solutions is available from the Intel Developer site.

For more information about the RadiSys design building blocks, visit the products area of the RadiSys site.

Author Bio

Mark Summers, a technical marketing engineer in the Embedded Intel Architecture Division, has been with Intel since 1999. Before joining Intel, Mark was with Motorola for 16 years. Mark has been granted 14 U.S. patents. He earned a B.S. in mechanical engineering from Arizona State University.

Desktop

Flash the Intel® BIOS With Confidence

Patrick Lang
Senior Product Marketing Engineer
Desktop Board System Software
Intel Corporation

Overview

Updating a system BIOS (basic input/output system)—known as "flashing the BIOS"—can resolve many common hardware-related issues in a computer system. Unfortunately, it's also a great way to render a system completely unbootable should something go wrong during the update. Consequently, many information technology (IT) professionals and technical support organizations view the risk/reward ratio of BIOS updates as excessive, and shy away from them.

The Intel® BIOS, exclusive to Intel® desktop boards, completely transforms the once risky task of updating a system BIOS by pairing superior BIOS protection with an easy-to-use BIOS update utility. These make BIOS updates safe, easy, and convenient. Intel patented fault-tolerant technology prevents corruption of an Intel BIOS during a flash update, and Intel patented digital signature technology provides virus protection by preventing unauthorized access to the BIOS flash part. Add to this Intel® Express BIOS Update, which makes updating an Intel BIOS as simple and easy as installing a new Windows* device driver, and you've got a winning combination.

Why Update a BIOS?

Computers exist in a changing ecosystem where new technologies and hardware devices continually enter the market and upset the current state of affairs. Computer systems continually face issues from existing devices that contain bugs and emerging devices that employ new technologies. BIOS updates can resolve many such issues, which can be limited to the device in question, or be severe enough to cause complete system instability.

One common example involves systems not waking up from standby mode properly. In a typical scenario, a system works fine until a new riser card or peripheral is added and problems surface in getting the system to wake up. Since the BIOS plays a key role in preparing a system for standby mode, a BIOS update can resolve such issues.

Another example involves new hard drives entering the market at the time of this writing. The latest drives are built with a new addressing architecture, allowing them to be larger than ever before. Six months ago, very few, if any, BIOSes supported this architecture. With digital video editing and other storage-intensive operations becoming more popular, users are adding these huge hard drives to many existing systems. A BIOS update is the easiest and most expedient way to get them to work in existing systems.

Risky Business: Problems Updating a Typical BIOS

It's easy to corrupt many non-Intel BIOSes during an update, especially if a power loss or shutdown occurs during the flash update. Most motherboard vendors release BIOS updates on a regular basis. While they usually provide detailed instructions for the upgrade procedures, the chance of BIOS corruption is a known risk.

Along with a power failure, other common problems include an improperly coded BIOS file, a bad flash routine, or a variety of other issues. For example, some manufacturers open the door to BIOS corruption by allowing a board to be updated with the incorrect BIOS. This can render the board unbootable.

Intel® Flash Architecture Avoids Typical Problems

The Intel BIOS addresses these risks with increased security through two exclusive features: Intel patented fault tolerant flash technology and Intel patented digital signature technology.

With fault-tolerant flash technology in place, you cannot corrupt an Intel BIOS due to complications during a flash update. If anything goes wrong during the flash, the BIOS automatically recovers and is re-flashed. The user doesn't have to open the chassis or touch anything on the board.

At the same time, digital signature technology improves virus protection by allowing only a signed BIOS image to be written to the flash part. In addition, the Intel BIOS allows itself to be updated only by the proper Intel update, preventing the wrong BIOS from being flashed onto any Intel desktop board.

How Fault-Tolerant Flashing Works

The Intel BIOS stores a backup boot block for safekeeping. If any problem occurs with the default boot block when the system POSTs (runs the Power On Self Test—the process BIOS performs to initialize the system before the operating system takes control), the BIOS automatically uses the backup boot block to force the system into BIOS recovery mode.

Advantages of Fault-Tolerant Flashing

Fault-tolerant flashing offers significant advantages, including:

- Removing the need for a backup BIOS—an expensive solution seen on some boards that involves a second (backup) flash part.
- Removing the need for "boot-sector protection" through jumpers or a setup option. (Boot-sector protection is a way of preventing writing to the boot block, but it also prevents the boot sector from being safely updated.)
- Reducing original equipment manufacturers' (OEMs) board returns due to BIOS corruption during a flash update.
- Enabling OEMs to confidently provide protected BIOS updates to their customers.

How Digital Signature Technology Works

A unique digital signature is embedded in every Intel BIOS. If the code attempting to write to the Intel BIOS does not provide a BIOS image with the correct digital signature, the existing BIOS will not allow reprogramming of the flash part. Digital signature technology also utilizes security features in the Intel® chipset, providing several rings of protection around the BIOS.

Digital Signature Technology Stands Up to Viruses

Intel patented digital signature technology enhances protection against attacks from viruses and rogue code by preventing unauthorized writes to the flash part. The Intel BIOS has weathered many virus attacks unscathed, underscoring the strength of its virus protection.

In one telling example, not a single case of corruption by the Chernobyl (or CIH) virus was reported on an Intel desktop board. This virus, first discovered in June 1998 and activated in April 1999, caused an estimated $250 million in damage to over 600,000 computers worldwide (Source: Victor Latona, "CIH: One Year Later," ZDNET.com, December 15, 2000).

Get Intel® Express BIOS Update Today

Intel Express BIOS Update allows for convenient and easy BIOS updates right from the Windows* OS, without the confusing procedures common to traditional BIOS updates. It's literally as easy as 1-2-3:
1: Download or e-mail the BIOS update (it's small enough to fit on a floppy disk).
2: Double-click the executable from Windows.
3: Click through an easy InstallShield* wizard.

Paired with the extraordinary protection built into the Intel BIOS, Intel Express BIOS Update opens up new avenues in the world of BIOS updates:

- End users—even technophobes—can easily and safely update their system BIOS.
- IT can push BIOS updates across a network with system-management software.
- Support organizations can resolve system issues efficiently by providing BIOS updates.

All BIOS updates for Intel desktop boards are available in a downloadable Intel Express BIOS Update package from the Intel Developer Web site.

Summary

There's no need to avoid updating the system BIOS on an Intel desktop board. Intel patented fault-tolerant flash technology and Intel patented digital signature technology can eliminate the risk of complications during a BIOS update and prevent viruses and other damaging code from overwriting the Intel BIOS.

Now OEMs, customer-support organizations, developers, and end users can all take advantage of easy-to-install BIOS updates that simplify life, reduce support costs, and solve many hardware/peripheral problems.

More Info

For more information on Intel desktop boards and other resources for developers, check out the Intel Developer site.

Author Bio

Patrick Lang has been with Intel for two years, and is currently a senior product marketing engineer in the Desktop Platform Solutions Division. He has also worked in product management for the Intel® WebOutfitter℠ Service as well as in software development and software technical marketing at other companies. Patrick received his B.S. in civil engineering from the University of South Carolina and his M.S. in hydrology and water resources from the University of Arizona.

Initiatives & Technologies

InfiniBand* Architecture Development Is Moving Forward

Jim Pappas
Director of Initiative Marketing
Enterprise Platform Group
Intel Corporation

Overview

The Internet is creating an increased demand for server computer I/O subsystem performance, scalability, reliability, and flexibility. An industry-wide shift to a fabric-based I/O architecture is required. The InfiniBand* Trade Association has developed an industry specification for a channel-based, switched-fabric architecture that provides a scalable performance range of 2.5 Gb/s to 30 Gb/s per link, meeting the needs of systems ranging from entry-level to high-end enterprise. With wide support from the computing industry's leading companies, InfiniBand Architecture represents the industry's choice for I/O technologies that will keep pace with the demands of the Internet age.

In this article you will learn about the history of InfiniBand architecture, its current status, and its immediate future. You also will learn about InfiniBand performance data and fabric demonstrations presented at the US Intel Developer Forum Conference, Spring 2002.

Why InfiniBand Architecture?

The InfiniBand Architecture is ideally suited to address the demands generated by the rapid growth of the Internet and the convergence of data and telecommunications (voice, data, video, and storage) on the Internet. It is a key technology for server clusters as well as I/O for remote storage and networking within the data center.

Initially, InfiniBand fabrics will be used to connect servers with remote storage and networking devices, and other servers. InfiniBand architecture will enable systems to balance I/O capacity and server capacity so IT can add servers without having to add I/O adapters—to connect multiple servers without having to add I/O connectivity if it’s not at full capacity. They can share that I/O capacity across many servers, providing savings along with simplification.

Other benefits of using InfiniBand architecture include higher performance, lower latency, easier sharing of data, built-in security, and quality of service. InfiniBand architecture can reduce total cost of ownership by boosting data-center reliability and scalability. It addresses reliability and cost issues by creating multiple redundant paths between nodes (reducing hardware that needs to be purchased). It also moves from the load-and-store-based communications methods used by shared local bus I/O to a more reliable message-passing approach.

InfiniBand architecture addresses scalability in two ways. First, the I/O fabric itself is designed to scale without encountering the latencies that some traditional data interconnect technologies experience as workload increases. Second, the physical modularity of InfiniBand technology will eliminate the need for customers to buy excess capacity up-front in anticipation of future growth. Instead, they will be able to buy what they need at the outset and “pay as they grow” to add capacity without affecting operations or installed systems.

Where InfiniBand Architecture Has Been?

InfiniBand architecture has been promoted since its inception two years ago by the InfiniBand Trade Association. The Association is led by seven computing-industry leaders—Compaq, Dell, Hewlett-Packard, IBM, Intel, Microsoft, and Sun Microsystems—dedicated to developing a new common I/O specification to deliver a channel-based, switched-fabric technology that the entire industry can adopt. InfiniBand architecture is based on the collective research, knowledge, and experience of these industry leaders.

The InfiniBand Trade Association has completed specifications, and trade association members have delivered samples and initiated first-generation demonstrations of databases on InfiniBand fabrics.

Where InfiniBand Architecture Is Now?

Demonstrations at the Intel Developer Forum Conference Spring 2002 (IDF) featured real-life data center applications running on an InfiniBand fabric. These demonstrations showed the progress Intel and other leaders have made in developing solution stacks on InfiniBand architecture.

These demonstrations featured the next level of InfiniBand solutions involving IBM DB2, Oracle 9*i* RAC, SASEnterpriseMiner, and Biokey software. More than 20 vendors demonstrated interoperability connecting InfiniBand architecture with Ethernet and Fibre Channel and highlighting multiple InfiniBand wire protocols running over a single fabric.

It's crucial to have so many vendors demonstrating interoperability because the environment in which InfiniBand architecture operates is characterized by many different products and technologies. Industry leaders working together to deliver diverse InfiniBand applications enables progress on delivering the promise of InfiniBand fabric connectivity to IT departments.

Intel showed a clustering demonstration for high-performance computing featuring crash-test simulation software provided by the NCSA (National Center for Supercomputing Applications). This demonstration highlighted the combination of the MPI protocol with InfiniBand fabric connectivity and measure performance gains versus traditional HPC (High Performance Computing) interconnect technologies. The performance and reliability attributes of InfiniBand architecture make it the interconnect of choice for future HPC solutions.

What's Next for InfiniBand Architecture?

The InfiniBand architecture industry is focused on rolling out InfiniBand fabric connectivity to data centers this year. On the near horizon is publication of performance data measuring the comparative performance attributes of InfiniBand fabrics to traditional interconnect technologies. Early data was made available at IDF, with more comprehensive industry figures becoming available as application stacks mature.

Intel is calling for the industry to work on a common set of APIs for InfiniBand architecture. Common APIs will enable cross-platform interoperability of InfiniBand products across the industry.

IT managers are starting to take acute interest in how InfiniBand architecture can solve their data center problems. Intel is focused on working with InfiniBand ecosystem players to deliver early IT-ready solutions.

Summary

As I/O demands continue to increase in the data center, InfiniBand architecture will address these demands by providing a unified I/O fabric for all data center connectivity. The flexibility of InfiniBand product development offers solutions for all segments of the industry. The intent of InfiniBand architecture is to be always at the leading edge of commercially implementable technology for high-performance communications.

Intel is a key player in leading the way to InfiniBand solutions. Intel is developing and supporting the application of innovative technologies that advance the promise of InfiniBand architecture, and is providing the means for these technologies to be tested and developed into InfiniBand architecture-enabled product solutions.

From helping to guide the industry initiatives that enable InfiniBand architecture to become a reality to creating the Intel Interoperability Lab and Intel InfiniBand Software Optimization Lab for product testing, Intel is committed to ensuring successful, industry-wide adoption of InfiniBand architecture.

More Info

See more on InfiniBand Architecture at the Intel Developer Site and the InfiniBand Trade Association site.

Author Bio

Jim Pappas works with the industry on the development of InfiniBand products and promotes the standard through the InfiniBand Trade Association. He has served on the boards of numerous technology initiatives, most notably as a founding member of the PCI Special Interest Group and as founding chairperson of the USB Implementers Forum. Previously, Jim was director of Technology Initiatives in Intel Desktop Products Group. During his 18 years in the computer industry, he has successfully advanced desktop technologies including AGP, DVD, IEEE 1394, Instantly Available PC, and USB. Jim holds eight U.S. patents in the areas of computer graphics and microprocessor technologies. His B.S.E.E. is from the University of Massachusetts in Amherst.

IPv6 Router Solution Using Intel® Internet Exchange Architecture

Harsha Hegde
Senior Software Engineer
Intel Labs
Intel Corporation

Dylan Larson
Business Development Manager
Intel Labs
Intel Corporation

Overview

Telecommunications equipment manufacturers looking to deploy Internet Protocol version 6 (IPv6) functionality seek avenues to reduce engineering effort and time-to-market. By leveraging innovative technologies based on the Intel® Internet Exchange Architecture (Intel® IXA) and a wide choice of standards-based network processing solutions from the third-party developer community, they can achieve rapid deployment of next-generation IPv6 networking and communications services, maximizing return on investment.

Market requirements driving IPv6 are gaining momentum. IPv6 satisfies the need to provide expanded addressing capabilities to serve the growing global Internet community, a well-documented limitation of today's predominant Internet Protocol version 4 (IPv4).

Temporary fixes devised to alleviate IPv4 addressing limitations such as private addressing, Network Address Translation (NAT), application layer gateways (ALGs), and other workarounds, have resulted in complex configuration requirements and have compromised routing efficiency along the Internet backbone. The overhead they've created is increasingly cumbersome and costly to provide. Flow label capability and the need for better support for real-time delivery of data—quality of service (QoS)—are also leading to an increase in IPv6 adoption. Improvements such as simplified header format, greatly enhanced security, and better support for extensions and options, are further driving IPv6 acceptance.

With many standard Internet applications beginning to support IPv6, many key technology leaders, such as Cisco, IBM, and Microsoft, are releasing IPv6 support for their products. The Yankee Group, a market analysis firm, believes the deployment of IPv6 is not a matter of if but when, and advises its clients to begin developing a migration strategy. (Kerravala, Zeus. "IPv6: Addressing the Needs of the Future," E-networks & Broadband Access, *FLASH*—Vol. 2, No. 11, Sept. 2001: The Yankee Group, Boston, MA.) To meet this market requirement, developers must prepare to make the crucial transition from IPv4 to IPv6.

The technology solution described in this article provides a layered software approach built on emerging standards that allows designers to use off-the-shelf hardware and software components from a variety of sources to shorten time-to-market. The solution also introduces a powerful programmable network packet processing architecture that allows designers to deliver wire-speed performance at a fraction of the cost and engineering effort of traditional approaches.

Technical Details

To quickly deliver reliable and high-performance access, edge, and core routing and switching solutions to market, designers need best-of-class components from a variety of sources that can be integrated without the investment of lengthy, custom engineering effort.

A router solution that meets these requirements is shown in Figure 1.

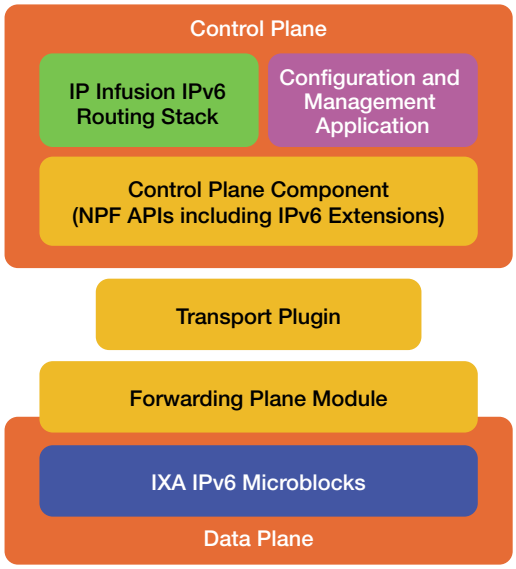

Figure 1. Innovative, modular IPv6 Router.

Control plane and data plane separation. Traditionally, network platforms such as routers were built as integrated devices. Routing protocols and management applications ran on a main processor, and ASICs were used to provide fast packet processing. A new approach separates the network device into two elements: the control plane and data plane (or forwarding). In this approach, the routing protocols and management applications reside on the control plane that is running on a general-purpose processor. The data plane is implemented using programmable network processing units (NPUs), which offers the flexibility to add new services without changes in packet forwarding hardware.

When the control plane and data plane are separated, two communications elements are needed:

- A standardized interface to manage communication between the data plane and the control plane
- An interconnect-independent protocol for communication between the two planes

The Network Processing Forum (NPF) and Application Program Interfaces (APIs)

Formed in February 2001, the NPF was created to encourage the growth and effective use of network processing technology. One NPF goal is to develop standard interfaces that integrate control plane software, such as routing and switching protocol stacks, with network processing components. By building upon NPF-proposed standards and specifications, integration between control plane and data plane components is simplified and provides designers with broader choices for development.

These current and future APIs from the NPF enable control plane/data plane communication:

- *IPv4 Forwarding API*—Used for adding and deleting routing/address resolution protocol (ARP) entries in the IPv4 forwarding/ARP table in the data plane
- *Interfaces API*—Used for configuring properties of interfaces in data plane
- *Packet Handler API*—Used for sending and receiving data packets
- *Multiprotocol Label Switching (MPLS) Configuration API*—Used for setting up and tearing down MPLS label switching protocols
- *IntServ API and DiffServ API*—Used for configuring Integrated Services and Differentiated Services Quality of Service mechanisms

In anticipation of NPF developments, Intel is also collaborating in the definition of the following IPv6 APIs:

- *IPv6 Forwarding API*—Used for adding and deleting routing and next neighbor entries in the data plane
- *IPv4/IPv6 Interoperability API*—Used for configuring several IPv6 transition mechanisms, including automatic/configured tunneling and 6-to-4 transition mechanisms

Forwarding and Control Element Separation (ForCES) Protocol

ForCES is an Internet Engineering Task Force (IETF) working group that is defining a messaging protocol for communication between the control plane and the data plane. As part of this effort, it is defining the logical functions that operate between the two planes and the data elements that are carried as part of the operations between them. The messaging protocol supports a variety of interconnect methods and supports communication between a control plane element and multiple data plane elements.

IP Infusion* IPv6 Routing Stack. IP Infusion is a member company of the NPF, collaborating on the standardization of IPv4 control plane to data plane interface specifications. IP Infusion has contributed a series of draft API specifications to the NPF to facilitate control plane-to-data plane communications for routing and switching. IP Infusion is a leader in IPv6 with its routing solution deployed in hundreds of nodes around the world. IP Infusion's routings stacks includes support for the following IPv6 routing protocols:

- RIPng
- OSPF version 3
- BGP-4+

Intel® Control Plane Platform Development Kit (PDK)

Intel provides prototypical implementations of emerging NPF control plane APIs and the IETF ForCES interconnect protocol in its Control Plane PDK. The kit allows faster integration of networking software with data plane silicon, thus providing faster time-to-market. Intel CP PDK architecture is modular, as shown in Figure 2.

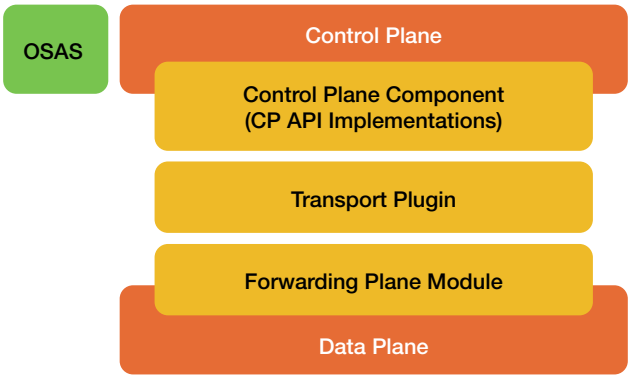

Figure 2. Intel Control Plane Platform Development Kit (PDK).

The CP PDK consists of three main components.

- *Control Plane Component*—Exposes ratified and proposed NPF APIs and hides hardware details. It supports control and management of multiple data planes.

- *Transport Plugin*—Provides communication between the control plane and data plane by implementing the ForCES protocol.
- *Forwarding Plane Module*—Translates ForCES messages to the data plane-specific interface. It exposes hardware capabilities to the control plane and forwards control and exception packets to control plane.

Intel's CP PDK is extensible and is written using an Operating System Abstraction Service (OSAS) so that it can be easily ported to other platforms.

Intel® IXA

Intel Internet Exchange Architecture (IXA) is a packet processing architecture that provides a foundation for software portability across multiple generations of network processors. It is based on programmable microengine technology, the Intel® XScale™ microarchitecture and the Intel IXA hardware abstraction layer.

A large Intel IXA third-party developer community provides numerous standards-based, high-performance Intel IXA communications building blocks. These building blocks support faster time-to-market, greater design flexibility, and extended time-in-market for next-generation networking solutions. Intel IXA enables rapid deployment of differentiated, reliable, and intelligent services while maximizing return on investment.

Intel® IXA IPv6 Microblocks

Microblocks are software components for processing packets within the Intel IXA programming framework. Microblocks run on the microengines of an Internet Exchange Processor (IXP) and use hardware features for fast packet processing. Examples include microblocks for IPv4 forwarding and Network Address Translation (NAT). For IPv6 packet processing, the following microblocks are required.

- *IPv6 forwarding microblock*—This microblock processes IPv6 packets coming from a Receive microblock. If the destination address is a local address, the packet is passed to IPv6 stack driver to be sent to the control plane. If the destination address is a remote address, the next hop is determined by a forwarding table lookup. The packet is then handed over to Transmit microblock to be sent out on a particular port. Control packets, such as Internet control message protocol (ICMP) packets, are passed to a stack driver, which in turns sends them to the control plane.
- *IPv6/IPv4 interoperability microblock*—This microblock supports IPv6 transition mechanisms that permit co-existence of IPv6 and IPv4. As several IPv6 transition mechanisms require IPv6 in IPv4 tunneling, this microblock provides encapsulating IPv6 packets in IPv4 headers to support transition mechanisms such as automatic tunneling, configured tunneling, and IPv6-to-IPv4 transition mechanisms.

Figure 3 shows IPv6 and IPv4 microblocks working together.

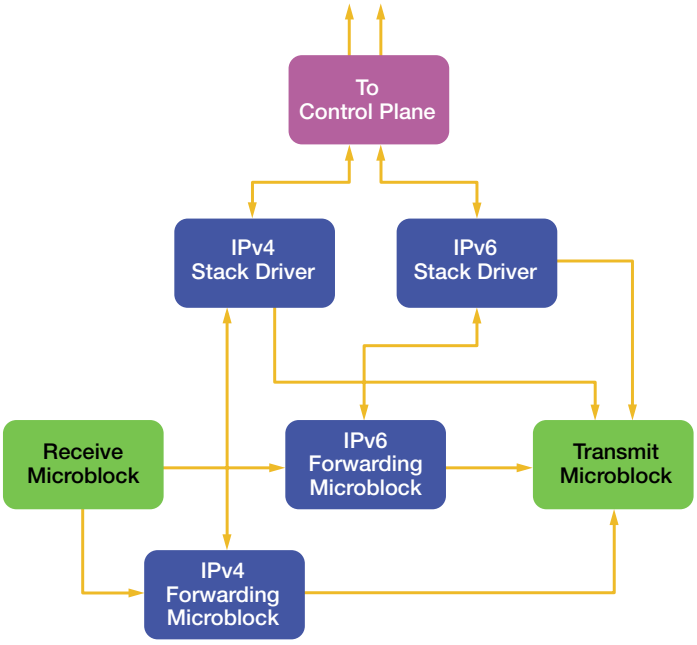

Figure 3. Intel IXA IPv6 and IPv4 Microblocks.

The IPv6 forwarding and IPv6 transition microblocks greatly simplify the engineering effort needed to support next-generation IPv6 routing solutions.

Summary

Control plane and data plane separation based on standard interfaces allows innovation at multiple levels in the stack. Intel IXA provides a highly programmable data plane using network processors. This allows network equipment providers to:

- Choose the best software and hardware components from a broad selection to build network devices
- Reduce engineering time and optimize time-to-market
- Deploy new network services quickly and efficiently

Developers can achieve rapid deployment of next-generation IPv6 networking solutions by taking advantage of Intel's innovative programmable network processors and the Intel Internet Exchange Architecture (Intel IXA).

More Info

For more information about IPv6, see the IP Version 6 Working Group Web site and the IPv6 Forum site.

For more information about the Network Processing Forum, visit the Forum's site.

More information about developments regarding the IETF Forwarding and Control Element Separation (ForCES) Protocol are available from an unofficial news source and archival repository for the IETF ForCES working Group.

Intel Internet Exchange Architecture resources can be found in the Networking and Communications area of the Intel site.

See the IP Infusion site for more information about their control plane and data plane separation technology.

Author Bios

Harsha Hegde is a senior software engineer in Intel Labs. He joined the company in 1997, and has worked on Quality of Service technologies such as RSVP and on Policy Based Network Management solutions, including the COPS protocol. Currently, Harsha is working on IPv6 and IPv4-to-IPv6 transition mechanisms for the Intel Internet Exchange Architecture (Intel IXA).

Dylan Larson, a business development manager in Intel Labs, is working on developing an Intel® cross-architecture strategy for IPv6. Previously, Dylan worked on ecosystem development for Intel Labs' Policy-Based Management and Quality of Service networking technologies. Before coming to Intel, Dylan was co-founder and director of product management for New Edge Networks. He is a member of the IPv6 Forum and has been a speaker at numerous broadband and ISP events.

Developing a Third Generation I/O Specification

Pat Correia
Initiative Marketing Manager
Intel Labs
Intel Corporation

Overview

As time progresses, the demands of emerging computing models will exceed the capabilities of the traditional PCI bus. There is a need for much greater internal system bandwidth, as technical innovations such as GHz CPU speeds, faster memory, higher-speed graphics, Gigabit LAN, 1394b, and others drive higher levels of system capabilities.

The Third Generation I/O (3GIO) architecture is a high-speed, general-purpose, serial I/O interconnect that will provide a unifying specification, consolidating the number of I/O interconnects within a platform, and eventually replacing the existing PCI, AGP and other derivatives with a single interconnect technology.

The Arapahoe Work Group is developing 3GIO specifications to help enable the development of products based on this technology. The work group consists of promoter companies, including Compaq, Dell, Hewlett-Packard, IBM, Intel, and Microsoft, as well as 61 key developer companies.

3GIO should meet the needs of multiple market segment platforms for both the short term and long term. This serial interconnect decreases interface pin count, resulting in a design that is more cost-effective, provides maximum bandwidth per pin, and is highly scalable to address future performance needs. It will be designed to serve as a long-term, general-purpose I/O interconnect to meet the requirements of desktop, mobile, server, communications, embedded, and other emerging applications.

Practical Benefits to System Designers

One of the main benefits that system developers should realize with 3GIO serial technology over a parallel bus architecture is the ability to deliver necessary I/O bandwidth in the fewest number of signal pins. System designers have the capability to scale (in both width and frequency) their designs to accommodate application needs while maintaining optimal cost-effectiveness. This capability provides system designers the flexibility of designing for performance while maintaining a competitive cost structure.

Another benefit and probably the most obvious to developers is that with 3GIO there are fewer lines to route on the motherboard than with existing parallel interfaces. Fewer lines means the connectors used for add-in cards are physically smaller. This inherent benefit allows designers more flexibility to either reduce cost or potentially add more features and/or capabilities to existing platforms as board space is freed up.

The 3GIO specification follows the established PCI software model for interoperability with the existing PCI paradigm; thus, 3GIO technology can run on today’s operating systems. This capability enables a smooth transition to new hardware while allowing software to evolve over time to take advantage of new 3GIO features.

The Hot Plug and Hot Swap features allow peripherals to be plugged or unplugged without powering the system down. System designers may also appreciate other design features of 3GIO technology, such as Quality of Service (predictable low latency transactions suitable for applications requiring isochronous data delivery, like video) and mechanisms to support embedded and communications applications.

Features to Support Transition of Upcoming Products

The 3GIO specifications also describe mechanical form factors to provide support for add-in modules and plug-and-play capability. In fact, there are two mechanical specifications: an Evolutionary Mechanical Specification and a Revolutionary Mechanical Specification.

As implied, the Evolutionary Mechanical Specification defines form factors for inclusion of 3GIO technology into existing chassis designs. The Revolutionary Mechanical Specification is intended to enable OEMs (original equipment manufacturers) and system manufacturers to develop flexible and innovative designs for system and cartridge form factors. This helps meet evolving market needs.

Some examples of these new system designs would split today's form factor into two smaller units to help make the system box more ergonomic or space efficient, or simply to have a design in a smaller form factor. Others would allow plug-in cartridge designs that would offer consumers an enhanced ease-of-use paradigm, as well as offering businesses and IT the ability to decrease total cost of ownership through easier serviceability and manageability.

Industry Working Together

Key developer companies were engaged in the second half of 2001 to bring cross-industry expertise to the development table through several versions of the draft specification.

At Intel Developer Forum Spring 2002 a release candidate version of the Third Generation I/O (3GIO) specification was announced. The release candidate specification, developed by the Arapahoe Work Group and made available to each of the key developers, signifies a key milestone toward delivering a draft 1.0 3GIO specification to the PCI-SIG at the end of the first quarter.

Transitioning the 3GIO Specification to the PCI-SIG

The Arapahoe Work Group will take the 3GIO architecture specification to its near final form and then transfer the document to the PCI-SIG. It is anticipated that the PCI-SIG will then carry the 3GIO specification forward.

The PCI-SIG has a large industry following from multiple market segments, and has achieved great success gaining industry adoption of I/O specifications. In addition, the PCI-SIG offers specification publication, compliance, and interoperability support, as well as technical support services.

Upon completion of the draft 1.0 specification transfer to the PCI-SIG for final review by PCI-SIG members, public release of the final 3GIO specification is anticipated by the second half of 2002. Products are expected to emerge in the marketplace by late 2003.

Summary

The 3GIO technology provides the industry with a long-term, general-purpose I/O interconnect to meet the requirements of desktop, mobile, server, communications, embedded, and other emerging applications.

This new architecture is being developed by key industry stakeholders in the Arapahoe Work Group and will be made available for broad industry adoption through the PCI-SIG.

The Arapahoe Work Group has just published a release candidate specification to its members, and plans to transfer the draft specification to the PCI-SIG by end of this quarter.

More Info

For more details on 3GIO technology and future plans, please visit the Intel Developer Web site and the PCI-SIG Web site.

Author Bio

Pat Correia is an initiative marketing manager in Intel Labs. He has worked at Intel for the past 10 years in chipset, processor, and initiative marketing. Pat received his B.S.E.E. degree from California State University at Fresno, and his M.B.A. from National University, San Diego, California.

Mobile

The Future of Mobile Computing—Four Vectors of Mobility

Ticky Thakkar
Director of Mobile Technology Development
Mobile Platform Group
Intel Corporation

Overview

When you think about what makes mobile computing valuable to PC users, four items stand out: the ability to run the most sophisticated and processing intensive applications, the ability to be connected to a network anytime, anywhere, the ability to stay mobile for extended periods of time, and the ability to carry and store the mobile PC easily. For designers and developers, and the manufacturers that support them, these items are being referred to as the "four vectors of mobility": high performance, seamless wireless connectivity, long battery life, and innovative form factors.

Intel is firmly committed to supporting designers and developers of mobile PCs to deliver the four vectors of mobility, through its own endeavors as well as in collaboration with other industry players. This article describes the efforts to deliver the four vectors of mobility: mobile processors designed and optimized to meet mobile user needs, fully integrated wireless networking technologies, technologies to make PC components more power efficient and batteries more long-lasting, and form factor technologies that enable even thinner and lighter mobile PCs.

Intel® Mobile Processor: High Performance and Low Power

Intel has been delivering high-performance, low-power microprocessors for the mobile PC segment for more than a decade. The Pentium® III processor-M is the latest member of Intel's mobile family of processors. The imminent introduction of the Pentium® 4 processor-M will bring an even higher level of performance while keeping power consumption low.

In 2003, the Banias processor will provide cutting-edge performance and even lower power consumption. Banias is the code name for the next-generation mobile processor that is being designed and optimized to meet the needs of the mobile users.

Mobile PCs based on Intel® mobile processors, such as the Pentium 4 processor-M and Banias, will benefit from performance advantages such as architectural features designed for advanced delivery of digital audio, video, photography, and enhanced 3D graphics, as well as accelerated graphics port (AGP) technology that facilitates the movement of large blocks of 3D data between the graphics controller and system memory. Even more remarkable is that Intel®-based mobile platforms will provide these performance advantages in an extremely power-efficient manner as a result of many state-of-the-art power management features. It is crucial in the mobile environment to deliver high performance while achieving low power consumption.

Seamless Wireless Connectivity That's also Practical

In terms of wireless connectivity, Intel's goal is to deliver a mobile PC that connects as seamlessly as a cell phone. This means that users can enjoy wireless or untethered connectivity to the network at all times as they move from floor to floor and from building to building within an enterprise, or from the home to the airport, or even from one city to another city on the opposite side of the globe.

Intel has put in place a number of research and product development projects to make seamless wireless connectivity not only possible but also practical. Researchers are developing technologies to address critical aspects of wireless communications, including seamless roaming, interoperability, ease of use, and wireless security.

On the product development front, Intel's direction on the wireless LAN silicon component is a major departure from the way wireless capabilities have been implemented in mobile PCs in the past. Before, such systems relied on an add-in card, most commonly a PCMCIA card. By integrating a wireless LAN solution, including an internal antenna, on the motherboard, Intel is eliminating the compatibility and interoperability issues that plague add-in card-based wireless technologies.

A related issue is the radio antenna or antennae that a mobile PC depends on for wireless networking. Implementing more than one antenna in a mobile PC has been problematic for developers. So Intel is working on a multimodal antenna that can be used for multiple frequencies. Similarly, Intel is addressing communications conflicts between proximate mobile PCs. In the past, mobile PCs within a few meters of one another could not simultaneously use different network protocols because of frequency conflicts.

Long Battery Life Makes All the Difference

Wireless connectivity and high performance won't mean a thing if mobile users cannot stay mobile because of limited battery life. That's why Intel is aggressively working to further reduce idle power consumption in processors and chipsets, licensing more power-efficient voltage regulator designs that support Intel® processors and motherboards, and streamlining platform and device power-state management.

Moreover, Intel is working closely with industry players to reduce the power consumption of other components in mobile PCs. Examples are power-efficient screen technologies, such as backlight technologies. In addition, Intel is a strong supporter of the Advanced Configuration and Power Interface (ACPI), which specifies the manner in which the operating system, motherboard, and peripheral devices communicate with one another regarding power usage.

Intel also recognizes that mobile users need a way to tell how much battery life is left at any given moment. In collaboration with Duracell, Intel created the Smart Battery System, an open battery and computer interface specification enabling battery power readings that are more than 98 percent accurate.

Innovative Form Factors: From Portability to True Mobility

While high performance, seamless connectivity, and long battery life are the key functional aspects in mobile computing, size and weight are also among the most important considerations when a user is shopping for a mobile PC. Through its advanced processor packaging technologies, Intel has long been a leader in enabling thinner and lighter mobile PCs. With the Pentium® II processor, for example, Intel introduced the industry's first small-footprint surface-mounted Pin Grid Array (PGA) package, and with the Pentium® III processor, the company introduced a thinner micro-Flip Chip Pin Grid Array (FCPGA) socketable package and an even thinner surface-mount micro-Flip Chip Ball Grid Array (FCBGA) package.

Today, Intel is focusing its packaging technology enhancements on improved heat-transfer capabilities and, in particular, on thermally enhanced IC packaging solutions. Such solutions use optimized micro fans and heat exchangers to maximize cooling capacity in a limited space. They also use fluid-filled heat pipes and low-resistance interface materials to ensure efficient heat removal from the processor while simultaneously keeping temperature low on the external surface of the PC.

Summary

To help designers and developers build successful mobile solutions for business users and consumers, Intel is working diligently to deliver the four vectors of mobility: high performance, seamless connectivity, long battery life, and innovative form factors.

More Info

To learn more about the four vectors of mobility, visit the business products and services section within the Intel Business Computing area of the Intel site.

Author Bio

Shreekant (Ticky) Thakkar is the director of Mobile Technology Development in Intel's Mobile Platforms Group. He has over 24 years of industry experience in development and planning. At Intel, Ticky established direction for the mobile notebook PC's transition from portable to wireless. He also managed a business unit that delivered smart proactive services over wired and wireless devices and led the teams that developed the Pentium® Pro MP (Xeon™) processor and Multimedia/Graphics Media extensions to the Pentium III and 4 processors. Before joining Intel, Ticky pioneered the development of Shared Memory Multiprocessors and Databases at Sequent Computer Systems.

Ticky holds a Ph.D. and M.S. in computer science/engineering from the University of Manchester and a B.S. in computer science/statistics from the University of Manchester. He holds or has pending applications for over 50 patents. Ticky has published numerous articles and has edited special editions of IEEE journals.

Networking and Communications

Intel Announces Three New Network Processor Families

Doug Davis
General Manager, Network Processor Division
Intel Communications Group
Intel Corporation

Overview

Line speeds are increasing in every segment of the network, and the ability to cost-effectively deliver a rich new mix of services is a key issue for service providers and carriers. To help the communications industry meet the demand, Intel is taking network processing technology to the next level.

Intel was joined by more than two dozen third-party developers at the Intel® Developer Forum (IDF) Conference Spring 2002, as it announced three new network processor families, including the first network processors designed to enable the flexible delivery of rich network services at line speeds up to 10 Gbps. The new Intel® Internet Exchange Architecture (Intel® IXA) network processors, including the Intel® IXP425, IXP2400, and IXP2800 network processor families, represent the first complete line of network processors optimized for applications from customer premises equipment (CPE) to the core of the network.

Intel first presented the requirement for three families of network processors at the Fall 2001 IDF Conference and presented details of the new Intel IXA network processor architecture at the Microprocessor Forum last October. This official announcement of Intel's new network processor families demonstrates the growing momentum from developers who need programmable line-rate performance to differentiate their designs while balancing performance and cost and meeting short time-to-market deadlines.

At the Spring 2002 IDF Conference, 26 companies shared the stage with Intel to announce product plans based on the new Intel IXA network processors, with solutions ranging from operating systems and tools to switch fabrics, security solutions, boards, and developer services.

Three New Families: CPE to Core

The new Intel IXA network processor families provide developers with the modular building blocks they need to cost-effectively meet the packet handling requirements imposed by faster line rates and new network services:

- The new *IXP425 network processor* family is optimized for customer premises equipment, including residential gateways, small/medium enterprise routers, cable modems, wireless access points, and remote DSL Access Multiplexers (DSLAMs).
- The new *IXP2400 network processor* family is designed for OC-48/2.5-Gbps network access and edge applications including: WAN multi-service switches, DSLAMs, cable modem termination system (CMTS) equipment, 2.5G and 3G wireless infrastructure base station controllers and gateways, Layer 4–7 switches (content-based load balancers and firewalls), VoIP gateways, multi-service access platforms, high-end routers, remote access concentrators, and Virtual Product Network (VPN) gateways.
- The new *IXP2800 network processor* family is designed for OC-192/10-Gbps network edge and core applications including routers, multi-service switches, 10-Gbps enterprise switches and routers for data centers, storage area networks (SAN) in addition to IPSec and VPN solutions, and wireless Infrastructure equipment.

Intel® IXA Technology Advances

Each of Intel's three new network processor families is optimized to meet a discrete set of performance and interface requirements. While the packet processing requirements differ from one market segment to the next, the new Intel IXA network processor families have some important architectural characteristics in common.

- *Microengine technology* enables the high-speed data-plane processing required for deep packet inspection and handling at fast line rates. Intel® IXP2400 and IXP2800 network processors implement a parallel processing architecture based on a subsystem of fully programmable, multi-threaded microengines, which are RISC processing elements designed for high-performance packet processing in the data plane. The IXP425 network processor is based on a similar parallel architecture with multiple network processor engines.

Together with the IXP2400 network processor, the IXP2800 is the first implementation of Intel's *Hyper Task Chaining* processing technology, an innovative approach to managing data-dependant operations among multiple parallel processing stages with low latency. Hyper Task Chaining enables a single stream packet/cell processing problem to be decomposed into multiple, sequential tasks that can be easily linked together to enable deep packet inspection via software pipelining at line rate.

- *Intel® XScale™ microarchitecture* provides integrated application processing in the control plane. All of the new network processor families implement an Intel XScale core for high-density processing. Intel XScale microarchitecture provides the highest ratio of performance to power consumption in the industry, with core speeds up to 1,000 MIPS and power consumption as low as 10 mW.
- *The Intel® IXA Software Portability Framework* facilitates code development and reusability. The modular programming framework enables code portability among network processor-based implementations, as well as future generations of Intel IXA network processors.

Development Environment

The new IXP2400 and IXP2800 network processors are complemented by an integrated software development environment that enables developers to rapidly create new applications and to migrate code to future generations of Intel® network processors.

The Intel® IXA Software Developers Kit (Intel® IXA SDK) 3.0 builds on Intel's popular IXA software development environment. It enables developers to evaluate and tune network processors to meet specific application requirements.

The Intel IXA SDK 3.0 includes:

- An enhanced, graphical, cycle-accurate simulation capability that enables demonstration and debugging of software, independent of hardware
- A high-level Microengine C compiler to simplify development
- Full support for choice of leading embedded operating systems, including Wind River VxWorks* and MontaVista* Linux*
- Intel® IXA Portability Framework and libraries to enable development of modular, portable code blocks and integration of third-party software for code reuse, longer product life and easier maintenance
- Suite of completed networking building blocks and example applications to speed development

CPE: Intel® IXP425 Network Processor

With its distributed processing architecture, including the Intel XScale core and three network processor engines, the IXP425 network processor provides robust application processing headroom and wire-speed packet handling performance for WAN and LAN networking. Single-chip integration of data, voice, and security functions saves the cost of implementing separate devices and supports easier integration with other hardware.

The IXP425 also supports a wide range of flexible broadband and LAN access solutions through its multiple integrated interfaces. The integrated UTOPIA 2 interface supports up to eight ADSL or G.SHDSL PHYs or one VDSL PHY. The PCI version 2.2 Host and Option interface enables direct connection with devices such as 802.11x chips, PCMCIA controllers, Ethernet MACs, and cable MAC/PHYs. The network processor also provides two high-speed serial interfaces, which can serve as high-speed ports for direct connection to T1/E1 framers or to industry-standard SLIC/CODECS.

The IXP425 network processor supports two widely adopted operating systems. The wide range of third-party development tools and the Intel® IXDP425 Network Processor Development Platform complement the network processor by enabling fast evaluation, performance testing, development, and prototyping.

Access and Edge: Intel® IXP2400 Network Processor

As line rates increase and as demand continues to grow for new network services, the intelligence present in today's edge and core equipment will move to access products, just as services now implemented in the network core are migrating to the edge. Flexible service provisioning, scalable performance, and support for multiple protocols are all important requirements of edge and access equipment.

The IXP2400 network processor is optimized to meet these needs by providing 2.5-Gbps packet forwarding capability and traffic management functions on a single chip. The IXP2400 implements a high-performance, low-power 32-bit

Intel XScale core for processing complex algorithms, route table maintenance and system level management, in addition to a parallel packet processing architecture with eight fully programmable multi-threaded microengines that support more than 5 giga-operations per second.

Core: Intel® IXP2800 Network Processor

Equipment in the network core must be able to deliver the highest level of packet handling performance and high reliability. The IXP2800 network processor delivers 10-Gbps packet forwarding and traffic management on a single chip. Using Hyper Task Chaining technology, the multiprocessing subsystem of the IXP2800 network processor provides the total processing capacity required for complex handling of packets/cells at wire speeds that traditionally required the design of dedicated high-speed ASICs.

In an OC-192/10-Gbps link, the network processor must be capable of performing necessary Layer-3 through Layer-7 applications on a cell or packet in approximately 35 nanoseconds and then transmit the cells/packets in the correct sequence, without data loss. The IXP2800 implements a 32-bit Intel XScale core with 16 fully programmable multi-threaded microengines with the capability of performing more than 25 giga-operations per second, sufficient to process 40-byte POS packets at a line rate of 10 Gbps.

Summary

At the Spring 2002 IDF Conference, Intel announced three new families of Intel IXA network processors that provide the benefits of fully programmable, flexible wire-speed processing from CPE to the network core.

The new generation of network processors demonstrates the advantages of Intel IXA:

- Intel® microengine technology for deep packet inspection and forwarding at line rates that have traditionally required the development of fixed-function ASICs (IXP425 network processors implement network processing elements in place of microengines)
- Intel XScale technology, delivering industry-leading MIPs/mW performance for application/control plane processing
- The Intel IXA Software Portability Framework, to enable fast code development and reuse across multiple products and product generations based on Intel IXA network processors.

During the announcement, Intel was joined by 26 third-party solutions vendors who announced plans for a spectrum of supporting hardware and software building blocks. These third-party organizations form part of the Intel IXA Developer Network, now with more than 80 member companies.

This is your invitation to join us, and use the power of Intel IXA network processors to take your communications applications to the next level.

More Info

Contact your Intel representative for detailed information on Intel IXA network processors, including product briefs, datasheets, white papers and other collateral.

Information is also available at Intel Developer Site.

For the latest information on Intel IXA solutions from Intel and more than 80 companies in the Intel IXA Developer Network, visit the Web site.

Author Bio

Doug Davis is the General Manager of the Network Processor Division, where he manages the IXA Components Organization and activities related to the Intel® IXA Developer Network. He started his 17-year career at Intel as a product engineer in the Military Division and went on to manufacturing and operations. Doug was the general manager responsible for Intel's embedded 80186, 80386, and 80486 product lines, later becoming the operations manager for what is now the Embedded Microcomputer Division. Doug holds a B.S.E.E. from New Mexico State University and an M.B.A. from Arizona State University.

Servers

High-density Architecture Meets Electrical and Thermal Challenges

Stephen Montgomery
Senior Systems Engineer
Intel Labs Enterprise Architecture Advanced Systems
Intel Corporation

Will Berry
Mechanical Engineer
Intel Labs Enterprise Architecture Advanced Systems
Intel Corporation

Ven Holalkere
Manager
Intel Labs Enterprise Architecture Advanced Systems
Intel Corporation

Tomm Aldridge
Director
Intel Labs Enterprise Architecture Advanced Systems
Intel Corporation

Overview

As the use of rack mount servers continues to increase, especially in the data center-based Web serving, business logic, and decision support industries, the need for two features has become clear.

- Increased density of server functions is needed to allow growth of the workload without similar growth in plant, property, or support equipment. Without this increase in density, staying on a value track for computing in the data center is difficult at best.
- Volume servers in the data center must become easier to use and maintain. As the number and density of rack mount servers increase, so does the interconnect complexity and the degree of difficulty in cable management. Servers are now hard to repair or replace, and they suffer from reliability issues as a result of damaged or missed wired connections. Further, proliferation of the thick cable "mat" blocks the much-needed airflow that cools the servers.

Intel has developed a completely new methodology for the architecture of high-density servers that delivers a solution to these needs. By integrating server electrical and thermal connection to the rack and using high-speed serial (HSS) signaling, server density can be increased while greatly simplifying the serviceability and cabling of the server. This development, known as the Integrated Rack Cooling and Docking Station (IRCDS) is targeted at InfiniBand* and Gigabit Ethernet connected servers in the dual-processor and quad-processor space.

This new methodology, which will be available to server developers in the near future, overcomes electrical interconnect and thermal challenges to provide the reliability and efficient space utilization that data centers and other high-density server installations require. Server hardware designers who design equipment for high-density data center applications should start planning to apply this vital new technology.

Data Center Challenges

Data centers are specialized facilities that house Web sites. They contain automated systems that constantly monitor server activity, Web traffic, and network performance, and they report even very slight irregularities so that engineers can spot potential problems. Data centers have grown rapidly in the past few years, and now they face critical infrastructure problems.

- *Thermal cooling capacity limits*. Most data center equipment relies on air circulation to provide cooling. Because air is not an especially efficient heat-exchange medium, large volumes of air must circulate to cool each piece of equipment. As the density of equipment in data centers increases to better utilize expensive real

estate, however, the amount of air available for cooling per server decreases. Further, the flow of air is constricted by blockages from cabling, which increases thermal problems.

- *Cost of downtime*. In the fast-paced, highly competitive world of Web-based business, one hour of downtime can cost organizations thousands, or even millions, of dollars an hour (Source: LifeKeeper* for Microsoft Windows NT product brief, 25 Feb 2002). To prevent downtime, reliable and redundant systems are deployed, but systems failures and maintenance tasks still need to be performed. Traditional rack-mounted equipment, where cables attach to the back of the system in the rack, requires a high degree of human intervention to swap out a system. The process is time-consuming and prone to human error, and it can result in even more downtime.

New technologies, such as blade servers, were created to address these problems. A thin, modular electronic circuit board containing one, two, or more microprocessors and memory, a blade server is intended for a single, dedicated application (such as serving Web pages). It can be easily inserted into a space-saving rack with many similar servers. Blade servers share a common high-speed bus and are designed to create less heat, thus saving energy costs as well as space. However, current cooling technologies limit the density and performance of blades.

A New Methodology–and Solution

To solve data center problems without sacrificing flexibility, Intel has devised the Integrated Rack Cooling and Docking Station (IRCDS), an innovative new methodology for the architecture of server systems. In this approach, the rack becomes an electrical and thermal interconnect system for the servers it hosts.

The Integrated Rack Cooling and Docking Station (IRCDS) comprises two components:

- Rack docking station: technology available now (could deploy in one year or less)
- Rack cooling station: technology under development (could deploy in two to three years)

IRCDS server technology is built around a flexible architecture that lets users install memory, processors, and storage input/output (I/O) interchangeably–combining the cost-saving features of blades with the power of traditional server technology. IRCDS methodology has two key features:

- *Electrical interconnection*. All connections to the server are made to the rack, allowing for servers that are completely front-accessed and hot swappable. As a server is inserted into the rack, a motherboard interconnection couples the server to the rack. This eliminates the need to get behind the rack to disconnect and reconnect cables whenever a server is swapped out, reducing swap time from about 30 minutes to about two.
- *Thermal interconnection*. Each server in the rack is also connected to a central closed-loop water cooling system. A connector on the back of each server connects the rack's cooling system to a condenser inside the server. Water circulates in a sealed loop from the condenser to the processor, where the heat from the processor vaporizes the circulating water. The vapor then cycles back to the condenser, where it is cooled and returned to a liquid state.

IRCDS technology offers major benefits:

- *Rapidly serviced or swapped computers will save data centers from downtime.*
- *Remote thermal management solves high-density space constraints and allows better use of a data center's physical infrastructure.*
- *Both thermal and electrical design flexibility are increased.*

Rack Docking Station

The Rack Docking Station provides a "pluggable" electrical interface for rack-mounted servers (see Figure 1). The technology consolidates multiple interfaces such as keyboard, mouse, server management, and high-speed serial input/output (I/O) technologies such as InfiniBand into one connector. A 1-inch by 5-inch adapter board with receiving connectors is installed into the supporting hardware at the rear of the rack. This provides the socket for the mating connectors on the server.

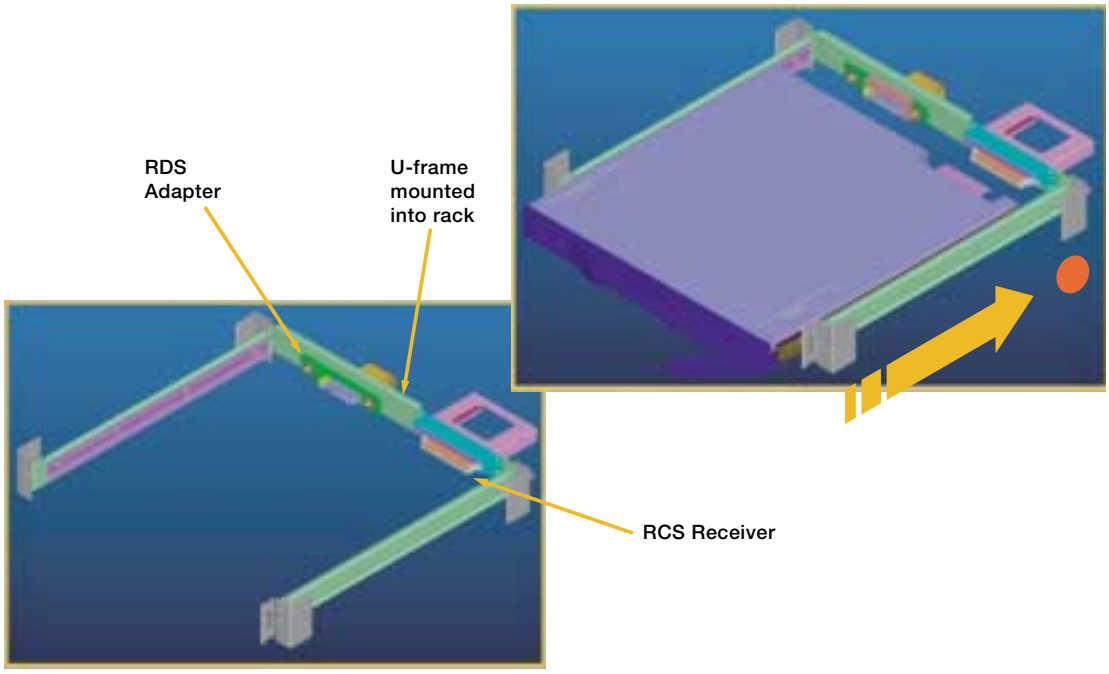

Figure 1. View of rack hardware and server insertion

Critical to the proper insertion of the interface is the ability of this adapter board to "float," to allow it to account for tolerance and other inconsistencies in the manufacturing of associated parts and still mate securely (see Figure 2). Silicon rubber bushings and shoulder screws allow the adapter board to float while chamfered edges on the alignment pins provide the "gathering" needed to ensure reliable, repeatable insertion.

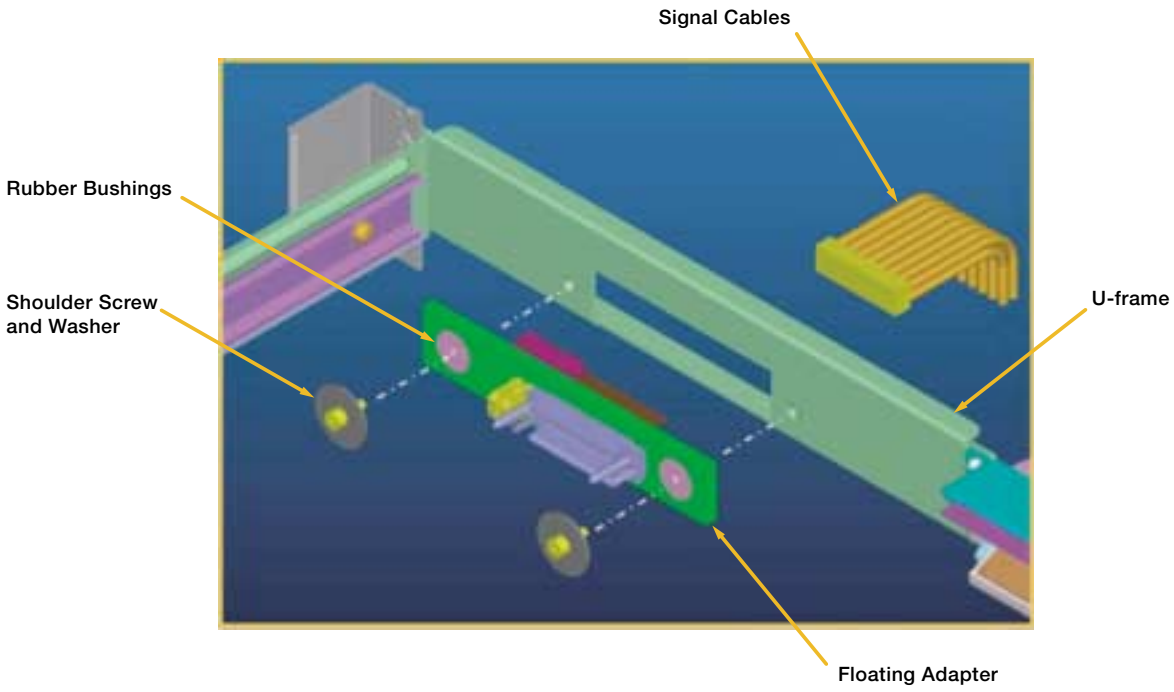

Figure 2. Exploded view of adapter board.

The principal benefit of this technology is increased serviceability for data center operators, who can quickly remove and replace servers from the front of a rack without having to replace interconnecting cables–a convenience that decreases service time and increases reliability of equipment swaps.

Rack Cooling Station

The Rack Cooling Station (RCS) is the thermal analog to the Rack Docking Station. RCS is a two-part system that allows for the removal of heat from small form-factor chassis to locations external to the chassis.

The first part, internal to the chassis, is based on highly reliable loop heat pipe technology already in use in critical applications such as avionics and satellites. It consists of an evaporator with a "wick" structure and a condenser connected together with a dual pipe system forming a loop, as shown in Figure 3.

The evaporator is attached to a heat source (a silicon package), and the condenser is connected to a "wedge" thermal docking mechanism. Inside the loop a small quantity of water is utilized as the working fluid. At the evaporator, the cooling fluid evaporates. With the fluid, vapors then traversing to the condenser via a pipe and are cooled to form droplets. These fluid droplets then traverse back to the evaporator, forming a two-phase fluid dynamic loop. This loop is entirely passive, and it requires no external work inputs to function.

The second part of the RCS is a receiver mechanism that is mounted external to the chassis. Heat is transported to an externally mounted cooling system via metal-to-metal contact between the wedge condenser and the receiver mechanism, as shown in Figures 3 and 4.

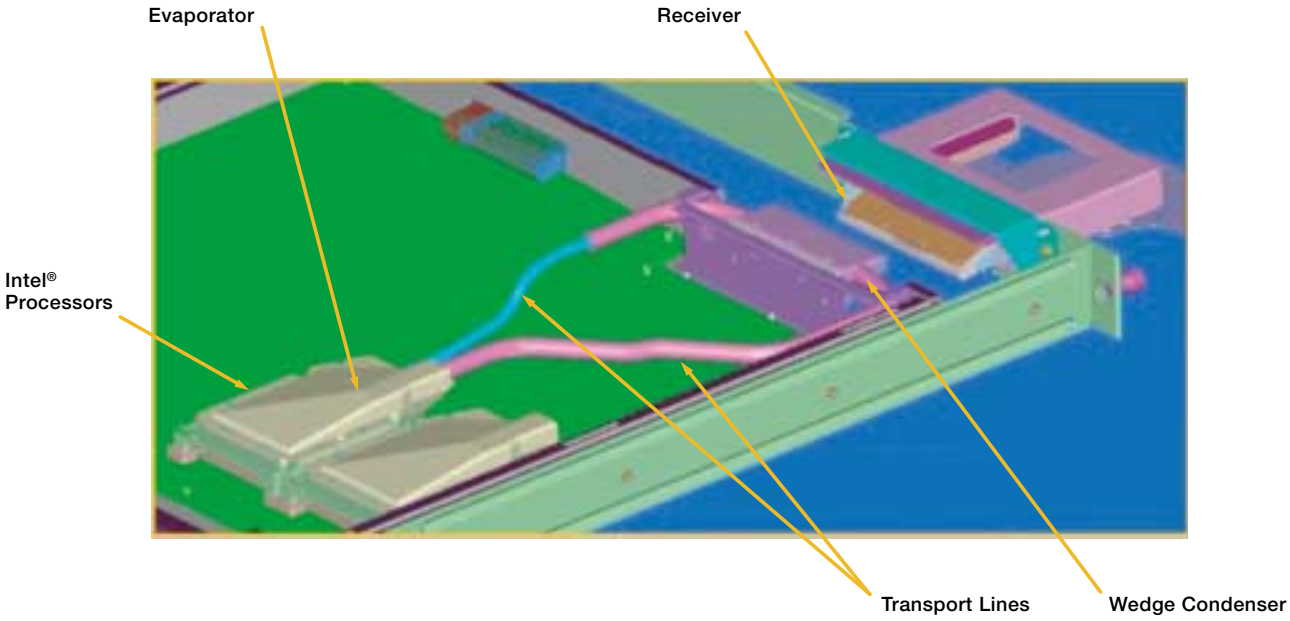

Figure 3. Rack Cooling Station Details

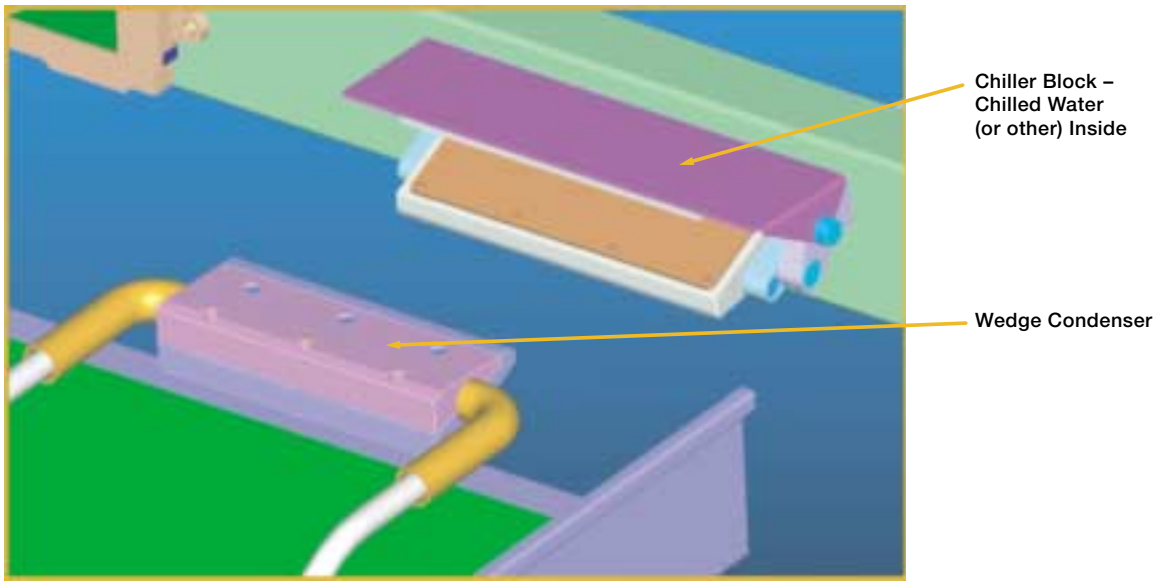

Figure 4. Receiver Details

The receiver mechanism is spring-loaded and supplies enough pressure to deliver a sufficient thermal interface to the wedge without requiring the use of a thermal interface material. External cooling fluid is supplied to the receiver mechanism in any manner the customer may desire. Some examples could include, but are not limited to, building chilled water, a refrigerated loop, or a pumped loop of a non-electrically conducting fluid.

Summary

Data centers that have grown in density to meet critical business needs face tough environmental and infrastructure issues. As the cost of real estate and downtime climb, technologies that increase reliability and reduce maintenance and downtime solve some of the challenges they face.

The IRCDS rapid swap technology is a completely new methodology for the architecture of high-density servers. It will save data centers from downtime and solve high-density space constraints to allow better use of a data center's building infrastructure.

More Info

For more information visit the Intel Developer site.

Author Bios

Stephen Montgomery is a senior systems engineer in the Enterprise Architecture Lab at Intel. He joined Intel in 2000 and has worked on the integrated rack cooling and docking station and on fouling of dense heat sinks. He has filed seven patents. Stephen holds M.S. and Ph.D. degrees in mechanical engineering from Purdue University.

Will Berry is a mechanical engineer in the Enterprise Architecture Lab at Intel. Since joining the company in 1997, he has worked on the dual Intel® Itanium™ processor-based concept server, dual Itanium processor-based chassis, dual Intel® Pentium® Xeon™ processor-based board and dual Pentium® II processor-based board. He has filed three patents. He holds a B.S. degree in mechanical engineering from Cornell University.

Ven Holalkere, manager of the Enterprise Architecture Lab, has been with Intel since 1999 and has 15 years of experience in the industry. His work history includes guiding thermal packaging and thermal research projects, developing advanced wire interconnect technology, and developing high-density servers. He holds an M.S. degree in mechanical engineering from West Virginia University.

Tomm Aldridge is the director and founder of the Enterprise Architecture Lab Advanced Systems Lab. Persistent "sockets" for Intel® Architecture-based enterprise servers are a key focus for the Lab, and IRCDS is aligned to that goal. Tomm joined Intel in 1995 with engineering, architectural, and management experience from his work in supercomputing and systems engineering with a variety of companies. He holds a B.S. in physics from the New Mexico Institute of Mining and Technology.

Wireless

Performance and Power Savings with New Applications Processors

David Rogers
Communications Marketing Manager
Handheld Computing Division
Intel Corporation

Tom Yemington
Product Marketing Manager
Handheld Computing Division
Intel Corporation

Overview

Developers seeking to successfully meet a range of price points in the wireless handheld market can look to the newly announced next-generation applications processors from Intel: the Intel® PXA250 and Intel® PXA210. As the first fully integrated products based on the Intel® XScale™ microarchitecture, these processors build on the success of the Intel® StrongARM* SA-1110 architecture with the performance and power savings essential for helping meet the needs of many full-featured handheld devices such as PDAs, multipurpose communicators, smart phones, and multimedia players.

Aiming squarely at the high end of the market with the highest processor speed available today (400 MHz) for handheld wireless devices, the Intel PXA250 processor delivers comprehensive integration, sophisticated multimedia capabilities, and superior power savings for the most advanced hardware designs and software applications. Targeting the value market, the Intel PXA210 processor also provides an advanced performance-to-power-consumption ratio but in a reduced-footprint, highly scalable solution ideal for small form-factor devices.

Successfully Balancing Performance and Power Savings

With its advanced 0.18µ CMOS process technology, Intel XScale technology gives the PXA250 and PXA210 processors a balance of performance and power consumption crucial for the multifunctional battery-powered devices increasingly in demand. With sizable on-chip memory caches and gated clocks for lower power dissipation, these processors provide one of the top MIPS/mW ratios in the industry. For example, running at 400 MHz, the PXA250 processor consumes less power than its predecessor did running at 200 MHz. Running at 200 MHz (as explained presently), the PXA250 (or PXA210) processor uses just a third of the power consumed by its predecessor at the same speed.

The Intel PXA250 and PXA210 processors also further fine-tune the performance/power-consumption balance by running at variable speeds: the PXA250 at 400 MHz, 300 MHz, and 200 MHz, and the PXA210 at 200 MHz and 133 MHz. Enabled by Intel® Media Processing Technology, this application-triggered power management capability scales up processor speed to meet the needs of a given application and then scales back down to conserve power when the application is finished running or no longer needs the additional MIPS.

Powerful, Scalable Multimedia Capabilities

To meet the demands of sophisticated multimedia applications such as Flash* animation and streaming video, the PXA250 and PXA210 processors include a 40-bit accumulator (“dual MAC”). This dual MAC (Media Access Controller) incorporates multimedia instructions implemented through Intel® Integrated Performance Primitives (Intel® IPPs). The Intel IPPs—which are low-level, cross-platform software algorithms for high-demand communications, signal processing, mathematics, and media functions—are designed to help reduce power consumption through efficient CPU execution so that developers can focus on value-added features and reduce their time-to-market.

To further assist developers in prototyping systems and building the applications of today and tomorrow, Intel is continually upgrading the Intel IPPs so they are optimized for the latest-generation Intel XScale microarchitecture. This means a developer building an MP3 player for a PDA or smart phone could write code to point to a downloadable Intel IPP-based MP3 decoder optimized for Intel Xscale microarchitecture. When targeting a future-generation Intel XScale core-based processor with the same application, the developer could leave the code unchanged, because the pointer would link to the latest Intel IPP decoder.

Highly Integrated and Expandable

So what exactly does an application developer or system integrator get with these next-generation processors based on Intel XScale technology? First, consider the Intel PXA250 processor.

The Intel PXA250 processor has its own LCD controller, an audio-codec interface, and a variety of communication ports, including USB with 16 end-points, a 1.84-MHz baseband (communications) interface, and Bluetooth* baseband interface at 920 Kbps. It features MMC/SD and PCMCIA/CF card support for expandable storage and I/O devices such as 802.11b and digital cameras. It also supports memory in both high-performance 3.3V/32-bit and lower-cost, lower-power-consumption 2.5V/16-bit versions.

Now consider the Intel PXA210 processor. The Intel PXA210 processor features the same capabilities as the PXA250 with the following exceptions: Instead of running at 400 MHz, 300 MHz, and 200 MHz, the PXA210 operates at 200 MHz and 133 MHz; it supports MMC/SD but not PCMCIA/CF cards; and it supports both 3.3V and 2.2V memory, but only in a 16-bit version. The Intel PXA210 also has a smaller footprint. Whereas the Intel PXA250 comes in a 17x17mm 256-pin PBGA package, the PXA210 comes in a 13x13mm 225-pin TPBGA package.

In sum, the Intel PXA250 processor is ideally suited for advanced PDAs and multipurpose communicators and the high-end multimedia applications designed for these devices, such as MP3 audio decode, MPEG-4 video decode, speech and handwriting recognition, and Java* interpreting. On the other hand, the Intel PXA210 processor is better suited for entry-level PDAs and smart phones and the mainstream, value applications designed for these devices.

A Rich Development Environment

The Intel PXA250 and PXA210 applications processors are key components of the Intel® Personal Internet Client Architecture (Intel® PCA), Intel's development blueprint for designing wireless handheld communication devices. A growing community of developers, called the Intel PCA Developer Network, offers wireless companies technical and marketing support for designing these devices and applications supporting Intel PCA. To date, more than 800 companies have joined the network and have access to its more than 400 hardware and software design tools.

Through the Intel PCA Developer Network, Intel has worked extensively with vendors to optimize key operating systems including Microsoft* Windows* CE.net, PocketPC* 2002, PalmOS*, SymbianOS*, and multiple versions of Linux*. This collaboration between Intel and the operating system vendors is designed to boost the performance of devices based on the Intel PXA250 or PXA210 processors running any of these operating systems.

Intel is also working with third-party compiler vendors to optimize their products for the processors, and a team at Intel is licensing advanced compiler technology to third parties. Finally, more than 200 independent software vendors are introducing versions of their popular multimedia software applications in support of these new Intel® processors.

On the hardware side, Intel is working with systems integrators and device manufacturers and helping designers build reference designs and development boards for products incorporating an Intel PXA250 or PXA210 processor.

Summary

The Intel PXA250 and PXA210 applications processors, the first fully integrated products based on the Intel XScale microarchitecture, can provide breakthrough performance and low power consumption for handheld wireless devices ranging from all-purpose phone / PDA / communicator / multimedia players to value cell phones and entry-level PDAs. Perhaps even more important, the scalability built into these processors can give developers and system integrators the headroom they need to compete and win, today and tomorrow, in this rapidly changing marketplace.

More Info

For more information on the Intel PXA250 and Intel PXA210 applications processors, availability, and pricing, go to the new Intel® Applications Processor Web site and the Intel® PCA Developer Network Web site.

Author Bios

David Rogers is communications marketing manager in the Handheld Computing Division at Intel Corporation, where he has been instrumental in taking the Intel XScale microarchitecture to market as well as promoting the next generation of processors based on the technology. Rogers holds a B.A. in political science from Austin College.

Tom Yemington is product marketing manager in the Handheld Computing Division at Intel Corporation. In his 11 years at the company, he has worked closely with Compaq Computer Corporation and Dell Computer Corporation. Yemington holds a B.S.E.E. and an M.B.A from the University of Texas.

—End of Intel Developer Update Magazine Issue 30—